SKA TELESCOPE MANAGER – PROJECT STATUS REPORT

L.R. Brederode*, SKA SA, Cape Town, South Africa
A. Marassi, INAF-OAT, Trieste, Italy
S. Riggi, INAF-OACT, Catania, Italy

Abstract

The Square Kilometre Array (SKA) will be the world's largest radio telescope once it is complete and will use hundreds of thousands of receivers, spanning Africa and Australia to survey the sky in unprecedented detail. The SKA will be ground breaking in many respects such as image resolution, sensitivity, survey speed, data processing and size to name a few. The SKA Telescope Manager Consortium is currently designing the SKA Phase 1 (SKA1) Telescope Manager Element that will orchestrate the SKA Observatory and associated telescopes. In this paper, we report on the current status of the SKA1 Telescope Manager pre-construction project, the development process and its high-level architecture.

INTRODUCTION

Telescope Manager (TM) is one of a number of elements within the scope of the overall SKA1 pre-construction project, where the TM element has three primary responsibilities [1]:

- Management of astronomical observations;
- Management of telescope hardware and software subsystems;
- Management of data to support system operations and all stakeholders;

The TM element can thus be seen as the nerve centre largely involved in controlling and monitoring the SKA telescopes by interacting with other elements, namely:

- Dish (DISH);
- Low-Frequency Aperture Array (LFAA);
- Central Signal Processor (CSP);
- Science Data Processor (SDP);
- Assembly, Integration and Verification (AIV);
- Infrastructure Australia and South Africa (INFRA-AUS & INFRA-SA);
- Signal and Data Transport (SaDT);

The TM pre-construction project was kicked-off in November 2013 and is primarily focused on performing requirements analysis, architectural design, interface definition and prototyping activities, with the aim to produce a data pack that enables procurement contracts to be placed during 2017. Once a contract has been awarded, the development of the TM will continue with low-level design, coding and testing, integration, qualification and installation related activities ultimately culminating with a transition of the TM into a fully operational state in 2022, as per current plans.

* rbrederode@ska.ac.za

At the time of writing, the TM pre-construction project had been in progress for roughly two years since kick-off. An international consortium, led by the National Centre for Radio Astrophysics (NCRA) in India, assumes responsibility for the delivery of the TM pre-construction phase data pack.

APPROACH AND STRATEGY

The development plan for the TM is based on a tailored set of lifecycle stages and associated processes derived from the IEEE 12207-2008 and ISO/IEC TR 24748-3 standards [2]. These standards have been chosen as a basis for the development plan given that the TM Element is predominantly comprised of software, and IEEE 12207-2008 is a recognised international standard that provides a comprehensive set of lifecycle processes, activities and tasks for *software* that is part of a larger system.

The technical processes are employed in an iterative manner at various levels within the TM product hierarchy, namely the element, sub-element and assemblies \ application levels [2]. Items at each level are subjected to a requirements and design review prior to a formal system-engineering baseline being formed. The baseline indicates that there is agreement in terms of what each item in the product hierarchy is required to do, what they will consist of, and how they will interact to achieve the required functionality. A critical design review (CDR) precedes the formation of the final design baseline during the pre-construction phase. Development of the various levels proceeds largely in parallel, however it is critical that a higher level be thoroughly reviewed and baselined before a lower level can be concluded.

The TM pre-construction project aims to deliver a detailed *system* design (element & sub-elements) and a high-level *software* design (applications). The construction phase project will commence with low-level software design.

A risk management plan sees to it that risks are identified, analysed, treated and monitored on a regular basis. Prototyping activities are aimed at mitigating many of the technical risks. Exploratory prototypes are used to prove design concepts, confirm technology choices, characterise interfaces and elicit requirements. These prototypes are not developed with reuse in mind; rather they will be discarded once they have served their purpose. This reduces their cost and time to delivery. The

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

THE LARGE SCALE EUROPEAN XFEL CONTROL SYSTEM: OVERVIEW AND STATUS OF THE COMMISSIONING

A. Aghababyan, R. Bacher, P. Bartkiewicz, T. Boeckmann, T. Bruns, M. R. Clausen, T. Delfs, P. Duval, L. Froehlich, W. Gerhardt, C. Gindler, J. Hatje, O. Hensler, J. M. Jäger, R. Kammering, S. Karstensen, H. Keller, V. Kocharyan, O. Korth, A. Labudda, T. Limberg, S. Meykopff, M. Moeller, J. Penning, A. Petrosyan, G. Petrosyan, L. Petrosyan, V. Petrosyan, P. Pototzki, K. Rehlich, S. Rettig-Labusga, H. Rickens, G. Schlesselmann, B. Schoeneburg, E. Sombrowski, M. Staack, C. Stechmann, J. Szczesny, M. Walla, J. Wilgen, T. Wilksen, H.-G. Wu, DESY, Hamburg, Germany
S. Abeghyan, A. Beckmann, D. Boukhelef, N. Coppola, S. G. Esenov, B. Fernandes, P. Gessler, G. Giambartolomei, S. Hauf, B. C. Heisen, S. Karabekyan, M. Kumar, L. Maia, A. Parenti, A. Silenzi, H. Sotoudi Namin, J. Szuba, M. Teichmann, J. Tolkiehn, K. Weger, J. Wiggins, K. Wrona, M. Yakopov, C. Youngman, European XFEL GmbH, Hamburg, Germany

Abstract

The European XFEL is a 3.4 km long X-ray Free Electron Laser in the final construction and commissioning phase in Hamburg. It will produce 27000 bunches per second at 17.5 GeV. Early 2015 a first electron beam was produced in the RF-photo-injector and the commissioning of consecutive sections will follow during this and next year. The huge number and variety of devices for the accelerator, beam line, experiment, cryogenic and facility systems pose a challenging control task. Multiple systems, including industrial solutions, must be interfaced to each other. The large number of bunches requires a tight time synchronization (down to picoseconds) and high performance data acquisition systems. Fast feedbacks from front-ends, the DAQs and online analysis system with a seamless integration of controls are essential for the accelerator and the initially 6 experimental end stations. It turns out that the European XFEL will be the first installation exceeding 2500 FPGA components in the MicroTCA form factor and will run one of the largest PROFIBUS networks. Many subsystem prototypes are already successfully in operation. An overview and status of the XFEL control system will be given.

INTRODUCTION

The European XFEL is a 3.4 km long X-ray Free Electron Laser (4th-generation light source) located in Hamburg, Germany. Its linear accelerator is based on superconducting RF-technology delivering 2700 electron bunches at energies of 10.5 GeV, 14 GeV, or 17.5 GeV with a bunch train repetition rate of 10 Hz. The undulator systems of the facility will produce spatially coherent photon pulses which are less than 80 fs long showing a peak brilliance of 10^{32}–10^{34} photons/s/mm^2/mrad^2/0.1% BW in the energy range from 0.26 to 29.2 at beam energies of 10.5 GeV, 14 GeV, or 17.5 GeV. Installation work in the accelerator and beam line sections is progressing. Technical commissioning and set-up with beam will be performed in stages in 2015 and 2016. First lasing should be possible end of 2016.

The control system of the European XFEL is an effort of various distinct groups at DESY and the European XFEL GmbH. This paper describes the concepts and features of the individual solutions for accelerator, undulator, photon beam line, and cryogenic and utility controls. It addresses the challenges to be faced, discusses the interoperability between the systems and reports on the status of the whole controls project.

ACCELERATOR CONTROLS

The accelerator control system provides an environment to supervise, control, synchronize and protect the technical equipment of the accelerator. It includes among other things a large scale front-end electronics infrastructure based on the novel MTCA.4 standard [1], a timing system, a machine protection system, server software and driver libraries, communication protocols, a finite-state machine tool to facilitate process automation, and a design tool for graphical user applications. In addition, it provides a central Data Acquisition System (DAQ) aggregating and storing shot-synchronized, bunch-resolved data from distributed sources and supporting various software-based feedback systems. The environment is completed by a code repository, a repository to keep the configuration data for generating graphical user panels, various Web services, and common IT infrastructure.

Conceptual Design

Primarily the conceptual design of the accelerator control system follows the DOOCS control system ansatz in place at the FLASH facility, itself a quasi-template for the European XFEL. DOOCS-based controls cover all aspects of equipment dealing with bunch-synchronous fast data taking and algorithms. In addition, slow control such as vacuum or magnet controls is also handled by TINE-compliant control system components. TINE is the principal control system of the PETRAIII complex (3rd-generation light source). Both DOOCS and TINE are full-featured control systems and have been recently tuned to interoperate seamlessly [2].

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

MAX IV LABORATORY, MILESTONES AND LESSONS LEARNED

V. Hardion, Y. Cerenius, F. Hennies, K. Larsson, J. Lidon-Simon, M. Sjöström, D.P. Spruce, MAX IV Laboratory, Lund, Sweden

Abstract

The MAX IV Laboratory is a new scientific research facility based on synchrotron light being built at Lund University, southern Sweden. The accelerator consists of one full energy linear accelerator providing two storage rings at 1.5 GeV & 3 GeV and a Short Pulse Facility. Additionally, more than 13 beamlines are planned to be built among which should be operational for the first users in 2016. The current status and approach of the control system is presented from its technical and organisational point of view, including the stakeholders, as well as the lessons learned from the commissioning as part of our continuous improvement for the future.

STATUS

Since the last ICALEPCS 2013 [1], the status of MAX IV Laboratory has moved toward the official operation milestones, the 21srt of June 2016 when the inauguration will take place. In the mean time the current MAX-lab (MAX I, II, III) will be decommissioned at the end of 2015.

MAX IV Commissioning

The Linac commissioning started in 2014Q2 and for the next 6 months it was focused primary on the 3 GeV storage ring injection. This milestone was reached in February 2014 despite the occurrence of several incident including the replacement of the Thermionic Gun.

The 3 GeV commissioning has just started in August 2015 after 10 months of equipment installation and Sub System Test (SST) but, already, several intermediate milestones have been passed:

- 14th August, 2015: Beam injected
- 26th August, 2015: Beam stored
- 8th October, 2015: Beam stacked at 4mA

The 1.5 GeV installation is going to be installed from the end of 2015 but already SOLARIS[28] which follow the same design has proven the concept, the 1st of October, by achieving 20mA at 1.5 GeV.

Beamlines

In the mean time, SPECIES, a Max II beamline and, FEMTOMAX, a Short Pulse Facility beamline received their first photon with a standard MAX IV control system beginning of 2014. Both are being used as test bench for the standard control system that will be used in the rest of the MAX IV beamlines [2]. BIOMAX, NANOMAX, BALDER, VERITAS, HIPPIE, all 5 3GeV beamlines, have successfully passed the site acceptance tests for their Front Ends and part of their optics.

KONTROLL & IT SUPPORT

The “Kontroll & IT Support” (KITS) group is responsible for providing the development and the support of the standard IT and electronics equipment, spread along the accelerators, the beamlines and also the administration of MAX IV. The group is part of the Science division and divided into several teams organised by business fields, responsible to organise the resource the most efficiently possible while keeping included into different projects as well as the other group of MAX IV, a so-called Hard Matrix organisation.

The KITS project management is based on Lean[3] where the focus is to bring the user value incrementally and as soon as possible without technical overhead. The Control System (CS) has been delivered following the Kanban[4] and the Scrum[5] methods since 2012[1], while adapting continuously the rules to follow the pace and the nature of a synchrotron project.

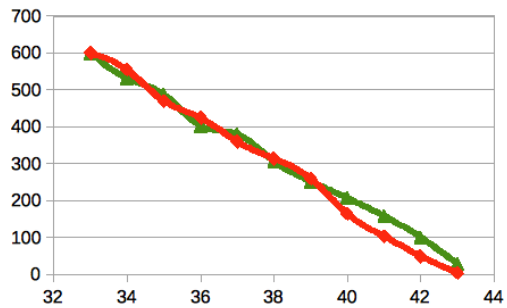


Figure 1: 3 GeV installation Burn down chart between sprint 32 and 44. Sprint estimation in red and done in green, in arbitrary unit

Network and System

The networks are based on TCP/IP v4 with a physical separation between Machine, Beamlines and offices. Each of them is divided by a set of virtual networks, usually corresponding to a subsystem, to minimise the influence of broadcast equipment. The layer 2 is based on Ethernet with 2 levels of switches in order to have a maximum network distance of 4.

All the accelerator and beamlines servers & IOC are virtual machines running by a VMWare hypervisor EsXi cluster of a redundant 36 cores and 384 GB memory. The standard OS is CentOS 6[6] as many providers support the RedHat drivers. All the Tango servers run on virtual machines apart for some specific case where a local control is needed. The definition of this infrastructure profits of the advantage that many equipment are controlled via Ethernet.

Ansible [7] is used for the configuration management of all the CS computers from the OS settings to the

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

THE CONSTRUCTION STATUS OF THE SuperKEKB CONTROL SYSTEM

M. Iwasaki#, A. Akiyama, K. Furukawa, H. Kaji, T. Naito, T. T. Nakamura, J. Odagiri, S. Sasaki,
KEK, Ibaraki, Japan
Y. Iitsuka, N. Yoshifuji, EJIT, Hitachi, Ibaraki, Japan
K. Asano, M. Hirose, KIS, Ibaraki, Japan
T. Aoyama, M. Fujita, T. Nakamura, K. Yoshii
Mitsubishi Electric System & Service Co. Ltd, Japan

Abstract

We have designed and upgraded the accelerator control system for SuperKEKB, the next generation B-factory experiment in Japan. The SuperKEKB control system is based on the features of the KEKB control system, while additional technologies have been implemented. In this paper, we report the construction status of the SuperKEKB accelerator control system.

INTRODUCTION

SuperKEKB is the upgrade of the KEKB asymmetric energy electron-positron collider for the next generation B-factory experiment in Japan [1]. The designed luminosity is to achieve a 40-times higher luminosity than the world record by KEKB. The KEKB operation finished in 2010 June to start the SuperKEKB accelerator construction. Currently the construction is in the final stage to prepare for the 1st operation in 2016 Feb. Here the accelerator operation will be performed in three stages, Phase 1, 2, and 3. The Phase 1 is for the vacuum scrubbing and basic machine tuning, and is the operation without the final focusing magnet system (QCS) nor the BelleII detector. After the Phase 1 operation, there will be Phase 2 and 3 operations with QCS and BelleII without (Phase 2) or with (Phase 3) the vertex detector.

The KEKB control system was based on EPICS (Experimental Physics and Industrial Control System) [2] at the equipment layer and scripting languages at the operation layer. The SuperKEKB control system continues to employ those features, while we implement additional technologies for the successful operation at such a high luminosity.

We have developed the interface modules to control thousands magnet power supplies, and introduce the EPICS embedded PLC, where EPICS runs on a CPU module. In the timing system, the new configuration for positron beams is required. The faster response abort trigger system has been developed. For the stable status monitoring, we have introduced the new alarm system based on CSS (Control System Studio) [3]. For Phase 2 and beyond, we have developed the new data archiving system based on CSS. The new VME-based FPGA module is also developed.

This paper describes the design and status of the SuperKEKB accelerator control system.

#masako.iwasaki@kek.jp

HARDWARE INTERFACE FOR SuperKEKB

For KEKB and SuperKEKB, EPICS is used as the main software to control the accelerator at the equipment layer. The EPICS framework consists of the Operator Interfaces (OPIs) and the Input/Output Controllers (IOCs). In KEKB, our system was constructed with the EPICS base version 3.13. For SuperKEKB, we're going to update to the version 3.14, however there still remain IOCs which are based on the EPICS base v3.13. For the operation layer, several script language of SAD script, python are used.

In KEKB, VME single board computers with VxWorks are mainly used as IOCs. For SuperKEKB, as well as the VME/VxWorks, the PLC (programmable logic controller) with a CPU module where Linux is running, and the Linux PC are also used as IOCs.

Figure 1 shows the picture of the PLC modules for the beam mask controller. In the CPU module (Yokogawa F3RP61), Linux is running, and we install EPICS into the CPU module to control the PLC. For SuperKEKB, the F3RP61 module is used to control the many devices of, for example, the vacuum system, LLRF, the automatic calibration subsystem of the large-type-magnet power supply, and beam collimators [4].

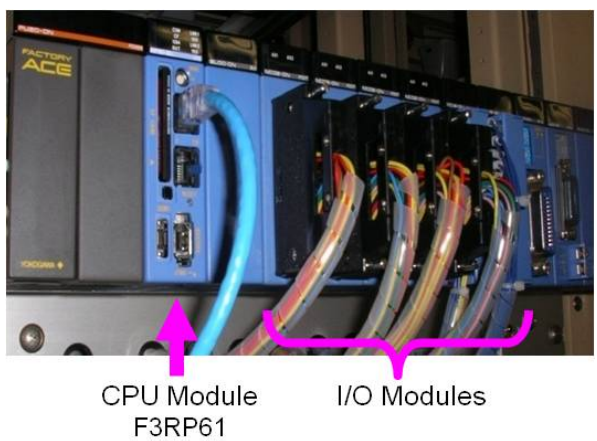


Figure 1: Picture of the PLC modules for the beam collimator controller.

There are many kinds of fieldbus in SuperKEKB accelerator components, such as Ethernet, GP-IB, serial, VXI/MXI for BPM, ARCNET for magnet power supply, and CAMAC. For the magnet-power-supply control, we have developed the PSICM (Power Supply Interface

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

COMPREHENSIVE FILL PATTERN CONTROL ENGINE: KEY TO TOP-UP OPERATION QUALITY*

T. Birke†, F. Falkenstern, R. Müller, A. Schälicke
Helmholtz-Zentrum Berlin für Materialien und Energie, Berlin, Germany

Abstract

At the light source BESSY II numerous experiments as well as machine development studies benefit from a very flexible and stable fill pattern. The fill pattern control engine is based on a state machine that controls the full chain from gun timing, linac pulse trains, injection and extraction elements as well as next shot predictions allowing triggering the next user experiment DAQ cycle. Architecture and interplay of the software components as well as implemented functionality with respect to hardware control, performance surveillance, reasoning of next actions and radiation safety requirements are described.

INTRODUCTION

After starting user operation in 1998, the 3rd generation light source BESSY II provided users with synchrotron radiation in decaying beam mode with three injections per day for 14 years. During this time BESSY II offered specific support for time-resolved experiments, pioneered low-α mode with coherent THz-radiation and ps-pulses and established the most advanced femtosecond slicing [1] endstation with 100 fs pulses. A purity controlled high current camshaft bunch in the ion clearing gap enables pump/probe experiments as well as pseudo singlebunch experiments at full intensity with a mechanical chopper.

Switching to top-up mode [2] extended these possibilities with a thermally stabilized machine and higher average intensity. Today, more pseudo single bunch experiments with reduced intensity are possible by using pulse picking by resonant excitation (PPRE) [3] with a selected bunch close to the end of the gap.

All these specific bunches have to be topped up with minimal variation to programmable intensities.

CONSTRAINTS

Radiation Safety

Based on analysis of facility properties and various malfunction scenarios, several restrictions and modifications have been set up to guarantee minimized losses and to make all sources of possible damages measurable and hence allow beamshutters to be kept open during injections:

- Injection efficiency must be >60% for every shot.
- Average injection efficiency must be >90% over 4 h-blocks.
- Maximum allowed injection frequency is 0.1 Hz.
- Maximum increase in stored current per injection shot: 2 mA.
- Minimum current in booster. For reliable and accurate efficiency measurement, a minimum current of 0.3 mA during an injection is mandatory.
- Minimum and maximum current of stored beam together with a corresponding minimum lifetime. To control and limit the normal average loss, a limit of max. 60 mA/h has been defined. The nominal limits are 200 mA-300 mA stored beam with a minimum lifetime of 5 h.

Top-Up Interlocks

Two separate systems are responsible for efficiency measurement and interlock. During top-up operation, both have to permanently check all relevant parameters for compliance with the top-up radiation protection restrictions and approve top-up operation to continue. The efficiency interlock systems are part of the radiation protection system.

Any violation of any of the restrictions will activate the top-up interlock system and pause or terminate top-up operation while beamshutters are free to open. Top-up injections are continued as soon as all top-up conditions are restored. This may happen automatically (e.g lifetime temporarily below limit) or may have to be accomplished by operations staff after closing beamshutters (e.g. last injection efficiency below 60%).

The fill pattern control engine operates in both top-up user-runs (beamshutters unlocked and free to open) and commissioning or machine development runs (beamshutters closed and locked). The latter is possible even without switching interlock systems to top-up mode. So the strict limitations of top-up mode are not enforced in this case and the top-up interlocks are inactive.

Injection Frequency and Timing

A proper compromise between minimizing variation in bunch currents and minimizing the number of disruptions due to injection kicker pulses had to be settled. The currently consolidated injection-scheme is, to fill a maximum of 5 bunches (reduced to e.g. 1 or 3 when filling dedicated single- or slicing-bunches) at every multibunch injection shot with the highest possible bunch currents. The decay phase between two injections currently ranges from 10 s to 200 s with an average of 120 s and every injection shot adds between 0.3 mA to 1.8 mA to the stored beam (ring current variation during top-up operation: ~0.5%).

* Work funded by BMBF and Land Berlin
† thomas.birke@helmholtz-berlin.de

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

PID_TUNE: A PID AUTOTUNING SOFTWARE TOOL ON UNICOS CPC

R. Martí, B. Bradu, E. Blanco Vinuela, CERN, Geneva, Switzerland
L. Frutos, R. Mazaeda, C. Prada, University of Valladolid, Valladolid, Spain

Abstract

PID (Proportional, Integral and Derivative) is the most used feedback control algorithm in the process control industry. Despite its age and thanks to its simplicity in terms of deployment and its efficiency in most of industrial processes, this technique still has a bright future. The major challenge in using PID control is to find the optimal set of parameters to tune the controller. This may be a complex task as it mostly depends on the dynamics of the process being controlled. In this paper a tool able to provide the engineers a set of PID parameters in an automated way is proposed. The tool offers auto-tuning methods, both in open and close loop, and others can be added as it is designed to be flexible. The tool is fully integrated in the framework UNICOS (CERN Standard Control framework) and can be used to tune multiple controllers at the same time, directly from the supervision layer.

INTRODUCTION

The largest majority of controllers found at the regulatory layer in the process industry are of PID (Proportional, Integral and Derivative) type. At CERN, around 4900 PID controllers are employed on the LHC cryogenic control system, 870 are employed on cooling and ventilation installations. In total, more than 8000 PID controllers are used at CERN in control systems based on the UNICOS (Unified Industrial Control System) framework [1]. All these controllers are implemented in more than 300 different PLCs (Programmable Logic Controller). Therefore, a tool to automatically tune the PID parameters would be a valuable asset in terms of controller robustness and hence in improving uptime and operation of the CERN industrial plants.

Although the PID controller has only three parameters, it is not straight forward, without a systematic procedure, to find suitable parameters. The main goal of this work is to provide an auto-tuning tool where different methods are available. The initial implemented methods were selected to cope with real industrial process scenarios. Two main features are searched: easiness of use, to allow operators to perform the task, and flexibility to extend the tool with new methods to cope with complex processes.

The article is organised as follows: first, the real implementation of the tool is depicted, second, the different tuning methods included in the tool are described including their suitability on different processes. Then, real industrial use cases with different auto-tuning methods employed are depicted along with a performance comparison between them. Finally, the paper concludes with the most important ideas and future work.

TOOL IMPLEMENTATION

The PID structure taken into consideration to design the auto-tuning tool is the ISA (International Society of Automation) version used in the UNICOS framework. The control law can be seen in the formula (1)

$$K_c \cdot (e(t) + \frac{1}{T_i} \cdot \int e(t)dt + T_d \cdot \frac{d}{dt}e(t)) \qquad (1)$$

The tool is fully integrated in the UNICOS framework at the supervision layer, the WinCC OA SCADA (Supervisory Control and Data Acquisition) from SIEMENS. It has been implemented in the WinCC OA control scripting language based on ANSI-C. The PID algorithm remains untouched in the PLC (control layer) and only its parameters are set once the tuning process is validated. The architecture and data flow is depicted in Figure 1.

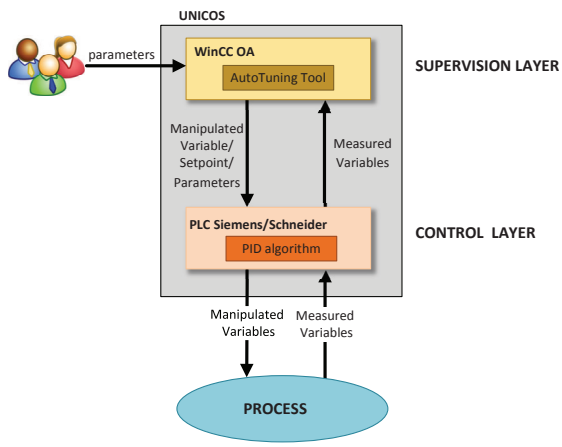


Figure 1: PID auto-tuning tool architecture.

The human interaction is made through dedicated interfaces embedded in the UNICOS controller faceplate. The user can select the method, configure its options, initiate the auto-tuning, evaluate the behaviour and ultimately set the new parameters to the controller. An example of such interface is shown in the Figure 2.

AUTOTUNING: BASIC CONCEPTS AND METHODS

The PID auto-tuning concept refers to the capability to automatically compute the parameters of a PID connected to the real plant. Many different approaches are available. This work focused on developing an auto-tuning tool, PID_TUNE, launched by the operators or the control engineers on their own initiative, performing the chosen tuning algorithm and

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

AUTOMATIC FEL OPTIMIZATION AT FERMI

G. Gaio, M. Lonza, Elettra Sincrotrone Trieste, Trieste, Italy

Abstract

FERMI is a seeded Free Electron Laser (FEL) located in Trieste, Italy. The machine setup and optimization is a non-trivial problem due to the high sensitivity of the FEL process to several machine parameters. In particular, the electron bunch trajectory and its spatial overlap with the seed laser beam represent one of the key aspects to optimize and then preserve during machine operation.

In order to ease the FEL tuning and to guarantee a long term stability of the photon beam, a software process integrated into the feedback systems performs automatic trajectory optimization of both the seed laser and the electron beams. The algorithm implementation, the results and the operational issues are presented.

INTRODUCTION

The FERMI machine setup is a non-trivial task. Although a deeper knowledge of FEL mechanisms in parallel with new diagnostic tools have reduced the time spent for machine tuning, new and more challenging requests coming from the users community have significantly increased the complexity of the accelerator setup and its operation during the experiments.

One of the priorities is the preservation of the FEL radiation quality in the long term. Over the years a number of feedback loops have been added in order to cope with subsystem instabilities [1]. However, even thermal drifts of only a few hundredths of degrees or requested changes of the machine working point tend to move the machine away from the optimum in terms of intensity, stability and spectral radiation purity.

In order to tune the machine for its best performance, a well established sequence of operations has to be carried out. At first, the photon-injector and seed laser systems have to be prepared, transversal and longitudinal electron beam dynamics in the linear accelerator (linac) have to be tuned, and then the FEL tuning process along one of the two undulator chains has to be performed. At the end, the photon beam is driven into the photon beam transport system and aligned in the experimental chamber. Before each users beam time one day is usually dedicated to machine preparation and tuning but, in case of non-standard experiments, up to five days have to be reserved for this task.

The tuning process which is performed by the physicists, in most of the cases consists of a number of manual or automatic scans of actuators versus machine variables, leaving to the physicist the final duty to set actuator to the optimum value. Instead, the optics matching, which imposes the design values of the Twiss functions to the electron beam [2], is completely automatic. There are two main reasons why the optics tuning is performed using the matching program: the first is that the problem is difficult to be managed (up to 6 dimensions in the linac optics). The second reason is that the model in terms of relations between quadrupole strengths and Twiss functions is well known. The availability of a system model reduces the optimization problem to a task that a proper algorithm quickly solves.

However, a system model cannot be always taken for granted. During the commissioning of a new system (e.g. a particle accelerator), the difference of the real system from the model could drive the model based optimization to frustrating results. In order to obtain a correct system model, software based simulators with adaptive and diagnostic capabilities have been developed and are currently supporting operations in several particle accelerators.

Although this kind of software could be applied to the entire FEL from the linac gun to the photon beam in the experimental chamber [3], its usage as an online tool has still to come. In particular, the multiplicity of the systems involved (photo-injector and seed laser beams, electron beam and FEL beam transport) and their nonlinearity have limited the adoption of these tools to only single well known subsystems, leaving to the manual intervention the rest of the machine preparation.

However, the manual tuning, which is a complex task, tends to be inefficient because the expertise of the personnel involved in machine operations is not homogeneous.

FEL PROCESS

The FERMI FEL design is based on an external seeding scheme. Two undulator chains, FEL-1 and FEL-2 operated one at a time, provide radiation in two different ranges, 100-20nm and 20nm-4nm respectively. In the modulator (first undulator) of both undulator chains, an external seeding laser in a tuneable range of 228-265 nm, overlaps in time and space the electron bunch provided by the linac. The resultant energy modulation of the electron bunch is then converted into charge density modulation by a chromatic dispersive section following the modulator. The electron bunch then passes through a series of undulators tuned in magnetic field so that coherent FEL radiation is emitted at one of the harmonics of the seed laser.

In FEL-2 the process is made in two stages. The radiation emitted by the first stage, that is similar to FEL-1, is used as seed in the second stage. In this way the seed laser wavelength is downshifted twice reaching shorter wavelengths.

The undulators and the seed laser are the most important devices involved in the tuning process. A software supervisor [4] taking into account the relations between the undulator magnetic strength, the electron beam energy and the seed laser wavelength, sets the

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

SYSTEM IDENTIFICATION AND ROBUST CONTROL FOR THE LNLS UVX FAST ORBIT FEEDBACK

D. O. Tavares#, LNLS, Campinas, Brazil
D. R. Grossi, São Paulo University, São Carlos, Brazil

Abstract

This paper describes the optimization work carried out to improve the performance of the LNLS UVX storage ring fast orbit feedback system. Black-box system identification techniques were applied to model the dynamic behavior of BPM electronics, orbit correctors, communication networks and vacuum chamber eddy currents. Robust control techniques were employed to analyze and optimize closed-loop performance and robustness.

INTRODUCTION

The LNLS storage ring FOFB system has been operating for users since March 2013. RMS stability integrated from 1 Hz to 500 Hz is 1% beam size in the horizontal plane and 6% for vertical in worst case (1% coupling is assumed). Disturbances originating on floor vibrations account for less than 0.2% RMS in horizontal plane and 2.5% RMS in vertical plane. The dominant disturbance source to be attenuated by the FOFB system is Elliptically Polarizing Undulator (EPU) gap/phase variations, while orbit disturbances caused by power supply ripple (up to 5% beam size, in vertical plane) should not be significantly amplified. The following sections describe the approach adopted to optimize the system performance and robustness.

HARDWARE DESCRIPTION

The LNLS FOFB system is a multiple-input multiple-output feedback control system comprising 48 sensors (24 beam position readings per transverse plane) and 42 actuators (18 horizontal orbit correctors + 24 vertical orbit correctors) running at 3.125 kHz update rate.

The data acquisition and control system hardware is composed of one central real-time controller (PXI-8108) and 12 data acquisition nodes (CompactRIO-9144 EtherCAT chassis) with analog and digital interface modules. Each acquisition/actuation node has an FPGA available for digital filtering, currently used for decimation (from 100 kS/s oversampling to 3.125 kS/s) and dynamic response compensation for each BPM reading and orbit correct current setpoint. Two daisy chain EtherCAT networks are used as deterministic network for sensor and actuator data distribution to/from the controller and provides synchronization between all nodes below 1 µs. Bergoz MX-BPM electronics are used for RF signals processing and delivers analog signals proportional to the beam position. In-house developed power supplies with analog interface and analog current regulator (Proportional-Integral type) are used to drive the orbit corrector magnets [1].

#daniel.tavares@lnls.br

SYSTEM IDENTIFICATION

Figure 1 depicts the general structure adopted for modeling the LNLS FOFB system using Laplace and Z transforms.

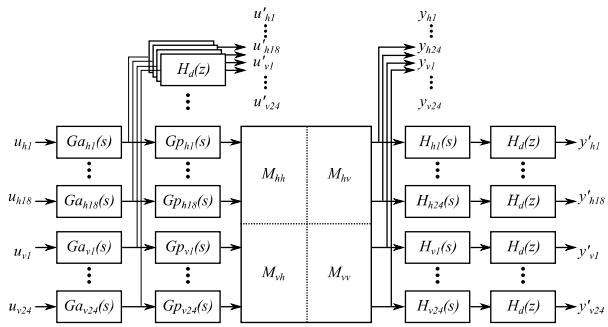


Figure 1: LNLS FOFB model structure.

The description of transfer functions and signals are given below:

- $Ga_{pj}(s)$: power supply + magnet impedance response of the j-th corrector of the p transverse plane (h = horizontal, v = vertical).
- $Gp_{pj}(s)$: magnet core + vacuum chamber magnetic field response (eddy currents) to coil current variations on the j-th corrector of the p plane.
- $H_{pi}(s)$: BPM electronics response of the i-th BPM of the p plane.
- $H_d(z)$: decimation filter response (discrete-time) at lower data rate (FOFB update rate) + communication network delay. Identical response for both BPM position reading and orbit corrector current reading.
- M_{p1p2}: static orbit response matrices ($p1$ plane BPMs and $p2$ plane orbit correctors). Horizontal and vertical plane matrices are represented by M_{hh} and M_{vv}, respectively. Crosstalk matrices (off-diagonal) are represented by M_{hv} and M_{vh}.
- u_{pj}: orbit corrector current setpoint (j-th corrector, p plane).
- u'_{pj}: orbit corrector current reading (j-th corrector, p plane).
- y_{pi}: actual beam position (i-th BPM, p plane).
- y'_{pi}: measured beam position (i-th BPM, p plane).

For black-box system identification the available ports for experiments comprise the discrete-time input signals u_{pj} and output signals u'_{pj} and y'_{pi}, with 320 µs sampling period. $H_d(z)$ transfer function is accurately known *a priori*, since it is fully determined by CIC decimation filters (factor 32 decimation rate, 2 sections and 1 differential delay) implemented in the nodes' FPGAs and one-cycle time delay due to the EtherCAT network data distribution.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

NSLS-II FAST ORBIT FEEDBACK SYSTEM *

Y. Tian#, K. Ha, L. Yu, W. Cheng, J. DeLong, L. Dalesio, BNL, Upton, NY 11973, USA
W. Levine, University of Maryland, College Park, MD 20742, USA

Abstract

This paper presents the NSLS-II fast orbit feedback (FOFB) system, including the architecture, the algorithm and the commissioning results. A two-tier communication architecture is used to distribute the 10kHz beam position data (BPM) around the storage ring. The FOFB calculation is carried out in field programmable gate arrays (FPGA). An individual eigenmode compensation algorithm is applied to allow different eigenmodes to have different compensation parameters. The system is used as a regular tool to maintain the beam stability at NSLS-II.

INTRODUCTION

NSLS-II is a third generation 3GeV storage ring with ultra-low emittance [1]. The low emittance requires a very stable electron beam orbit. Applying the common rule (beam stability < 10% of beam size), NSLS-II needs to hold submicron beam orbit stability. The stringent orbit stability requires the orbit feedback system to be able to suppress various noises from low frequency ground motion to high frequency mechanical vibration [2].

The NSLS2 BPM electronics generates high resolution turn-by-turn, fast acquisition (10kHz) and slow acquisition (10Hz) data. Fast acquisition data is used for global fast orbit feedback. To ensure sufficient bandwidth of the feedback loop, 10kHz data from all storage ring BPMs are distributed within a short period of time. A two-tier communication architecture is implemented to ensure the time budget for the FOFB system is satisfied. The first tier is the communication between local BPMs and the cell controller at each cell. The second tier is the communication between the cells around the whole storage ring.

The FOFB system is a typical multiple-input and multiple-output (MIMO) system. One common feature of the traditional singular value decomposition (SVD) based orbit feedback algorithm is that it applies the same controller dynamics (such as a PID controller) to all eigenmodes. Since each eigenmode has a different frequency response, it is desirable to apply a different controller for each of the eigenmodes and thus each eigenmode has different compensation in the frequency domain [3]. The challenge for such individual eigenmode compensation is that the feedback system needs to carry out much larger calculations within the time budget of the FOFB system. In this paper, we present the implementation and commissioning results of the FOFB system at NSLS-II.

*Work supported by DOE contract No: DE-AC02-98CH10886
#ytian@bnl.gov

TWO-TIER COMMUNICATION

One common task for a global orbit feedback system is to deliver BPM data to the calculation unit, and to send the calculated result, i.e. the corrector setpoints, to the power supplies. Many communication protocols are used in different labs.

To design the communication protocol at NSLS-II, we take advantage of the geographic locations of BPMs and correctors at NSLS-II. The NSLS-II storage ring is separated into 30 cells. There are 6 BPMs, 3 fast correctors, and 6 slow correctors in each cell. To effectively deliver the BPM and corrector data, we designed a hardware unit in each cell to collect local BPM data, deliver the BPM data around the ring, and to send corrector setpoints to the correctors. We call this unit the cell controller.

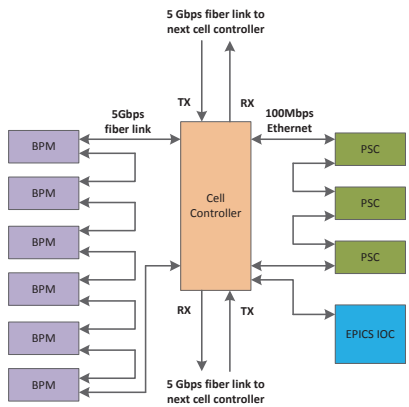


Figure 1: Two-tier communication.

By using cell controllers, two-tier communication is established, as shown in Figure 1. The first tier is the communication of the local data in one cell. In one cell, the cell controller and up to 12 BPMs are connected using fiber optic cables to form a ring structure. Currently, FOFB is only using 6 RF BPMs in the loop and it can be expanded to include 6 ID BPMs or X-ray BPMs. Each BPM sends its position data to its left and right neighbours, and passes through its neighbours' data. The data finally reach the cell controller and the cell controller saves the 6 local BPMs' data in its buffer. This communication protocol is called a serial device interface (SDI). On the other side, the corrector setpoints are sent from the cell controller to the local fast correctors via a similar SDI link. The only difference is that the power supply SDI link uses Ethernet PHY.

The second-tier communication is between the 30 cell controllers. The 30 cell controllers are connected using fiber optic cable to form a ring structure. Each cell

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

THE LASER MEGAJOULE FACILITY: THE COMPUTATIONAL SYSTEM PARC

S. Vermersch, J-P. Airiau, V Higonencq, J Pascal, X Julien, E Bordenave CEA, CESTA, F-33116 Le Barp, France
C. Lacombe, L. Burgy, J. Terral, A. Lemaire, Thales Services, Mérignac, France

Abstract

The Laser MegaJoule (LMJ) is a 176-beam laser facility, located at the CEA CESTA Laboratory near Bordeaux (France). It is designed to deliver about 1.4 MJ of energy to targets, for high energy density physics experiments, including fusion experiments. The assembly of the first line of amplification (8 beams) was achieved in October 2014.

A computational system, named PARC, has been developed and is under deployment to automate the laser setup process, and accurately predict laser energy and temporal shape. PARC integrates the simulation software MIRO. For each shot on the LMJ, PARC determines the characteristics of the injection laser system required to achieve the desired main laser output, checks the machine fine-tuning for equipment protection, determines the required diagnostic setup, and supplies post-shot data analysis and reporting.

This document gives the first results obtained with PARC. We also describe results obtained with the PARC demonstrator during the first experience on the LMJ facility.

PARC FUNCTIONALITIES

PARC must guarantee the laser performance according to a shot request. To achieve this goal, PARC is based on the simulation software MIRO which models each laser beam. PARC manages and updates one data model for each beam and validates them at the end of each shot.

Contribution in Experiment's Chronology

The LMJ monitoring system orchestrates the whole sequence of operations. It invokes PARC at each step of the sequence to deal with specific process.

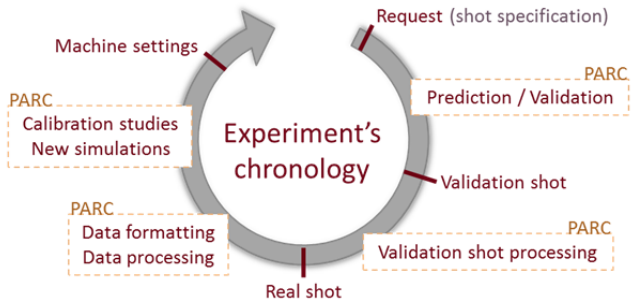


Figure 1: PARC in experiment's chronology.

The first step in the sequence is the production of a request that identifies all the desired laser output characteristics.

Predictive simulation is the first set of operations in the sequence chronology. It determines the LMJ equipment configuration and settings required to achieve the desired main laser output. The behaviour is performed by the software MIRO.

Validation is performed to check the ability of the machine to perform the shot in security and safety. Validation is also based on a simulation realized by MIRO.

Once the prediction and the validation are accepted, the sequence is composed of one or more validation shots to fine-tune data and equipment settings. A validation shot is like a real shot without the main amplifiers.

Data-processing of validation shot follows each validation shot. PARC performs data processing and provides parameters to check if the settings are compatible with the real shot. We also check that equipment's and diagnostics are operating normally.

The next step in the sequence is the real shot.

Data-processing of the real shot consists of supplying post-shot data analysis and reporting. Among other things, PARC compares the measurements to the request, to the results of simulations, performs new MIRO simulations from shots results to fine-tune calibration.

Calibration of computation models is performed at the end of each shot to correct the drift. The calibration uses the results and reports of several shots.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

LOW LEVEL RF CONTROL IMPLEMENTATION AND SIMULTANEOUS OPERATION OF TWO FEL UNDULATOR BEAMLINES AT FLASH

V.Ayvazyan*, S.Ackermann, J.Branlard, B.Faatz, M.Grecki, O.Hensler, S.Pfeiffer, H.Schlarb, C.Schmidt, M.Scholz, S.Schreiber, DESY, Hamburg, Germany
A.Piotrowski, Fast Logic, Lodz, Poland

Abstract

The Free-Electron Laser in Hamburg (FLASH) is a user facility delivering femtosecond short radiation pulses in the wavelength range between 4.2 and 52 nm using the SASE principle. The tests performed in the last few years have shown that two FLASH undulator beamlines can deliver FEL radiation simultaneously to users with a large variety of parameters such as radiation wavelength, pulse duration, intra-bunch spacing etc. FLASH has two injector lasers on the cathode of the gun to deliver different bunch trains with different charges, needed for different bunch lengths. Because the compression settings depend on the charge of bunches, the low level RF (LLRF) system needs to be able to supply different compression for both beamlines. The functionality of the controller has been extended to provide intra-pulse amplitude and phase changes while maintaining the RF field amplitude and the phase stability requirements. The RF parameter adjustment and tuning for RF gun and accelerating modules can be done independently for both laser systems. Having different amplitudes and phases within the RF pulse in several RF stations simultaneous lasing of both systems has been demonstrated.

INTRODUCTION

FLASH has been in operation as a user facility since summer 2005 [1]. In the mean time as a test facility, FLASH is in use for testing the superconducting accelerator technology for European XFEL [2] and ILC [3] projects. The first effort to operate accelerating modules at different gradients with alternating RF pulses was motivated by the ILC and European XFEL and initially has been demonstrated in year 2008 [4].

FLASH2 [5] is the second undulator beam line with variable gap built in a separate tunnel. It will make full use of the existing FLASH accelerator. Part of the bunch trains are kicked from the main beamline (FLASH1) into the new undulator beamline. In order to double the beamtime, users of both beamlines would need the 10 Hz repetition rate. A fast kicker in combination with a DC septum is used to deflect the beam into the second undulator line. In addition, the large variety in beam parameters should be available at both beamlines independently in order to ensure a maximum flexibility in planning of the beamtime. For this reason two cathode lasers are in use, each with its own bunch train. A variable delay between the two lasers within the RF pulse gun and accelerating modules ensures that users from both beamlines get their own set of parameters.

* valeri.ayvazyan@desy.de

In addition, the start time of the kicker pulse is shifted with respect to the start time of the laser pulse and the RF amplitude and phase of the gun and each of the modules can be tuned for optimal conditions for both users. During the initial tests it was shown that the RF system is able to handle the flexibility needed to compress the beam independently for FLASH1 and FLASH2 beamlines.

LOW LEVEL RF CONTROL FOR MULTI-BEAMLINES

RF Control Specifics for Multi-beamline Case

During the multi-beam line operation, within limits given by the beam delivery system, the bunch pattern and beam energy should be adjusted independently for each beam line, suggesting a time-sliced operation. For the FLASH2 beamline RF amplitude and phase changes within the pulse are required within a short time (less than 50μs). Different beam loadings are foreseen for different beamlines. Particularly the ability of gradient tuning of the last two RF stations is needed for wavelength scans for the FLASH1 beamline and the ability of phase tuning at injector for variation in compression at FLASH1 and FLASH2. From the operation point of view, the most important requirement is to have the ability of independent RF operational parameters adjustments for both beamlines.

LLRF System Overview

The layout of the FLASH facility including the new beamline is described in [6] and shown in Fig. 1. Its accelerator comprises a normal conducting RF gun, a first 8-cavity cryomodule, a third harmonic cryomodule with 4 cavities, a first bunch compressor, a second accelerating RF station (16 cavities), a second bunch compressor, and another two RF stations, with 16 cavities each. FLASH is operated in pulsed mode with repetition rate 10Hz. RF pulse duration is 1.3ms, 500μs for filling and 800μs flattop (beam acceleration). The RF power coming from 10MW klystrons is equally distributed to all cavities through the waveguide distribution system. The current stage of the LLRF control is implemented based on the MTCA.4 system [7]. The goal of the system is to control the accelerating gradient in amplitude and phase for each RF station based on vector sum control of cavity gradients [8]. The main components of the LLRF system are depicted in Fig. 2. For every cavity, the forward (PFWD), reflected (PREF) and transmitted signals or probes (PRB) are first down-converted to an intermediate frequency (IF) by the down-converters (uDWC) and then digitized

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

BAYESIAN RELIABILITY MODEL FOR BEAM PERMIT SYSTEM OF RHIC AT BNL*

P. Chitnis#, Stony Brook University, NY, USA
K. A. Brown, Brookhaven National Laboratory, NY, USA

Abstract

Bayesian Analysis provides a statistical framework for updating prior knowledge as observational evidence is acquired. It can handle complex and realistic models with flexibility. The Beam Permit System (BPS) of RHIC plays a key role in safeguarding against the faults occurring in the collider, hence directly impacts RHIC availability. Earlier a multistate reliability model [1] was developed to study the failure characteristics of the BPS that incorporated manufacturer and military handbook data. Over the course of its 15 years of operation, RHIC has brought forth operational failure data. This work aims towards the integration of earlier reliability calculations with operational failure data using Bayesian analysis. This paper discusses the Bayesian inference of the BPS reliability using a two-parameter Weibull survival model, with unknown scale and shape parameters. As the joint posterior distribution for Weibull with both parameters unknown is analytically intractable, the Markov Chain Monte Carlo methodology with Metropolis-Hastings algorithm is used to obtain the inference. Selection criteria for the Weibull distribution, prior density and hyperparameters are also discussed.

INTRODUCTION

The Beam Permit System (BPS) of Relativistic Heavy Ion Collider (RHIC) monitors the health of RHIC subsystems and takes active decisions regarding beam-abort and magnet power dump, upon a subsystem failure. The reliability of BPS thus directly impacts the RHIC downtime, and hence its availability. A Monte Carlo reliability model based on exponential competing risks was developed for BPS previously [1]. This model simulated the progression of basic component failures to system level catastrophic events. This work together with a quantitative fault tree analysis [2] helped characterize the failure rate and structural importance of each basic component of the BPS. RHIC has been operational for 15 years, and has gathered hardware failure data over this time period. This data represents the actual BPS failure attributes from a top level perspective. Bayesian analysis is a good candidate for combining these two information sources to get a combined inference about the system failure characteristics.

BAYESIAN PARADIGM

Bayesian statistics is branch of mathematics that deals with updating current knowledge about a system or process when new information is acquired. Statistical analysis follows two major approaches, namely frequentist and Bayesian. In the widely used frequentist approach, the probability distribution of an event is calculated by observing its occurrence over a large period of time, and the distribution parameters are assumed to be constant over time. In contrast, Bayesian approach keeps updating the probability distribution as new data arrives. The parameters of the distribution are treated as random variables that are modified according to the new information gathered. This becomes quite important when there are two sources of information about a system or process that indicate different results, and both sources hold significance to the inference.

The underlying framework for Bayesian analysis is Bayes theorem. Bayesian analysis involves the continuous form of the Bayes theorem [3], which is represented as

$$\pi(\theta|x) = \frac{L(\theta|x) \times \pi(\theta)}{f(x)}$$

The unknown parameter is θ, which defines the probability distribution of any process and is subject to change with the arrival of new information. Variable x is new source of information in the form of data observations. The term $\pi(\theta)$ is called as the prior distribution of θ, which can be elicited by using another parameter(s) called the *hyperparameter(s)*. $L(\theta|x)$ is the likelihood function for θ which is calculated by the gathering the new data. The term $\pi(\theta|x)$ is called the posterior distribution of θ, which is a combination of both prior and data likelihood function. $f(x)$ is the unconditional distribution of the variate x that acts as a normalizing factor in the equation. Because $f(x)$ is independent of θ:

$$\pi(\theta|x) \propto L(\theta|x) \times \pi(\theta) \quad (1)$$

This equation forms the foundation for Bayesian analysis discussed in this paper. We will discuss the selection of the prior distribution and data distribution (likelihood function) in subsequent sections.

PRELIMINARY ANALYSIS

Preliminary analysis is needed to find the suitability of Bayesian analysis to our problem and for choosing the distribution appropriate to the information sources we have. We analyze two the sources of information, the results from a Monte Carlo (MC) model [1] and the historical failure data obtained from the RHIC hardware maintenance records.

*Work supported by Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy.
#prachi.chitnis@stonybrook.edu

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

CONTINUOUS DELIVERY AT SOLEIL

G. Abeillé, A. Buteau, X. Elattaoui, S. Lê, Synchrotron SOLEIL, Gif-sur-Yvette, France
G. Boissinot, Zenika, Paris, France

Abstract

IT Department of Synchrotron SOLEIL [1] is structured along of a team of software developers responsible for the development and maintenance of all software from hardware controls up to supervision applications. With a very heterogonous development environment such as, several software languages, strongly coupled components and an increasing number of releases of the entire software stacks, it has become mandatory to standardize the entire development process through a "Continuous Delivery approach"; making it easy to release and deploy on time at any time. We achieved our objectives by building up a Continuous Delivery solution around two aspects, Deployment Pipeline [2] and DevOps [3].

A deployment pipeline is achievable by extensively automating all stages of the delivery process (the continuous integration of software, the binaries build and the integration tests). Another key point of Continuous Delivery is also a close collaboration between software developers and system administrators, often known as the DevOps movement.

This paper details the feedbacks on how we have adopted this Continuous Delivery approach, modifying our daily development team life and give an overview of the future steps.

SOLEIL CONCERNS&ANALYSIS

Synchrotron SOLEIL is a service center that delivers photons to external scientific users. The 24/7 operation of the Synchrotron SOLEIL facility relies on developed software by the "Control and Acquisition Software" team. This team has to regularly deliver changes requested by the business; it has become a challenge to perform it in the right timeframe. That is why the Continuous Integration and Delivery concepts have been studied. The main objective of these concepts is to automate every single process from the code building up to its operation in production.

Continuous Integration

Continuous Integration is a practice for detecting software defects at the earliest stage possible within the full development cycle to limit them via the automation of every step. Continuous Integration relies on key elements such as:

- *Version Control*: everything must be checked in to a single version control repository: code, tests, build and deployment scripts…
- *Automated Build*: Every project must be built and tested from the command line and must be runnable into the Continuous Integration environment.
- *Team agreement*: The commitment of the whole development team is required. Everyone must frequently check in small changes; and agree that the highest priority task is to fix any change that breaks the application.

Continuous Delivery

Continuous Delivery goes a step further. It assures to deliver a fully working and tested software in small incremental chunks to the production platform. It enables frequent release and deployment. All phases of a software product lifecycle from its coding up to production maintenance have to be automated (Fig. 1).

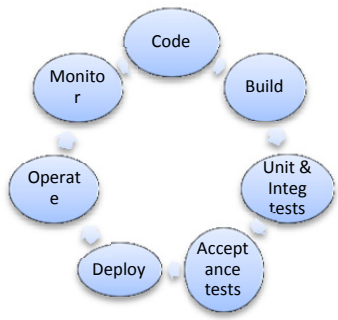


Figure 1: Continuous delivery.

DevOps

Soleil DevOps philosophy strengthens up a collaborative mode between the Development and Operations staff to achieve the same objectives: provide a full set of IT services meeting both business and IT requirements in terms of functionalities, responsiveness, and availability. DevOps is the mandatory journey to obtain a fully operational continuous delivery process.

Expected Benefits

The main expected benefits for SOLEIL of these methods are:

- Reduce all potential error-prone manual actions
- Reduce time to delivery; improve the reproducibility, repeatability and reliability of the delivery process.
- Federate the team around a common work process
- Focus on coding business requirements rather than on tooling
- Improve the synchronisation between change management and operation; reduce the number of incidents

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

SHOT RATE IMPROVEMENT STRIVE FOR THE NATIONAL IGNITION FACILITY (NIF)

G. Brunton#, G. Bowers, A. Conder, JM. Di Nicola, P. Di Nicola, M. Fedorov, B. Fishler, R. Fleming, G. Lau, D. Kalantar, D. Mathisen, V. Miller-Kamm, V. Pacheu, M. Paul, R. Reed, J. Rouse, R. Sanchez, M. Shaw, E. Stout, Y. Tang, S. Weaver, R. Wilson
Lawrence Livermore National Laboratory, Livermore, CA 94550 USA

Abstract

The National Ignition Facility (NIF) is the world's largest and most energetic laser experimental facility with 192 beams capable of delivering 1.8 megajoules of 500-terawatt ultraviolet laser energy. The energy, temperatures and pressures capable of being generated allow scientists the ability to generate conditions similar to the center of the sun and explore physics of planetary interiors, supernovae, black holes and thermonuclear burn. NIF has transitioned to a 24x7 operational facility and in the past year significant focus has been placed on increasing the volume of experimental shots capable of being conducted so as to satisfy the demand from the wide range of user groups. The goal for the current fiscal year is a shot rate of 300 (> 50% increase over the previous year), increasing to a sustainable rate of 400 the year after. The primary focus areas to achieve these increases are; making more shot time available, improvements in experiment scheduling, and reducing the duration of a shot cycle. This talk will discuss the control system improvements implemented and planned to reduce the shot cycle duration and the systematic approaches taken to identify and prioritize them.

INTRODUCTION

The National Ignition Facility (NIF) laser system [1] provides a scientific center for the study of inertial confinement fusion (ICF) and matter at extreme energy densities and pressures [2]. Each NIF experiment, or shot cycle, is managed by the Integrated Computer Control System (ICCS) [3], which uses a scalable software architecture running code on more than 2000 front end processors, embedded controllers and supervisory servers. The NIF control system operates laser and industrial controls hardware containing 66,000 control points to ensure that all of NIF's laser beams arrive at the target within 30 picoseconds of each other and are aligned to a pointing accuracy of less than 50 microns. Every NIF shot cycle [4] consists of approximately 1.6 million sequenced control point operations, such as beampath alignment, configuring diagnostics and arming triggers.

NIF has transitioned to a 24x7 operational facility and in the past year significant focus has been placed on increasing the volume of experimental shots capable of being conducted so as to satisfy the demand from its wide range of scientific user groups and maximizing the return of experimental data. The goal for the FY15 fiscal year was defined as a shot rate of 300 (> 50% increase over the previous year), increasing to a sustainable rate of 400 the year after. Figure 1 depicts the historical shot rate achieved on NIF and the goals going forward. Also indicated are the planned methods for achieving these goals:

- Formalizing a 24/5 shot week and increasing the number of weeks per year performing experiments (increased to 44 weeks)
- Improving the NIF shot scheduling paradigm by grouping similar shots to minimize the frequency of mid-week diagnostic reconfigurations
- Reducing the shot to shot durations. This includes both the shot cycle duration and the activities performed between each shot

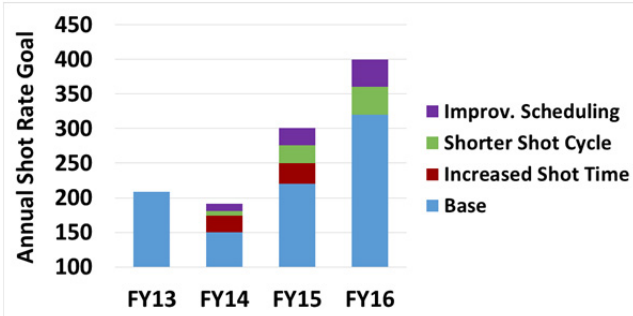


Figure 1: Historical and Planned NIF shot rates.

With both the increase in shot weeks and shot scheduling improvements commenced focus shifted to reducing the shot to shot durations. This work discusses the approach that was taken, a summary of the enhancements implemented and the results achieved.

ANALYSIS APPROACH

During each shot metrics are collected and archived that capture execution times of all phases and steps of the cycle. The mining of information from this vast data set was a critical component of the analysis and led to many of the identified improvements. To assist in rapidly visualizing these metrics the team leveraged the use of Splunk [5], a tool that we previously had extensively used

brunton2@llnl.gov

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

INTRODUCING THE SCRUM FRAMEWORK AS PART OF THE PRODUCT DEVELOPMENT STRATEGY FOR THE ALBA CONTROL SYSTEM

G. Cuni, F. Becheri, D. Fernandez-Carreiras, Z. Reszela, S. Rubio-Manrique, ALBA-CELLS Synchrotron, Cerdanyola del Vallès, Spain

Abstract

The Controls Section at the Synchrotron ALBA[1], produces and supports the software to operate the accelerators, the beamlines and the peripheral laboratories. It covers a wide range of disciplines like vacuum, motion, data acquisition and analysis, graphical interfaces, or archiving. Since the installation and commissioning phases, we have been producing the software solutions mostly in single-developer projects based on the personal criteria. This organization scheme allowed each control engineer to gain the expertise in particular areas by being the unit contact responsible to develop and deliver products. In order to enrich the designs and improve the quality of solutions we have grouped the engineers in teams. The hierarchy of the product backlogs represents the desired features and the known defects in a transparent way. Instead of planning the whole project upfront, we try to design the products incrementally and develop them in short iterations mitigating the risk of not satisfying the evolving user requirements. This paper describes the introduction of the Scrum framework as the product development strategy in a service oriented organization like the Computing Division at Alba.

INTRODUCTION

Alba is a 3rd generation Synchrotron Light facility located in Cerdanyola del Vallès, (Barcelona), commissioned in 2012 with seven operational soft and hard X-ray beamlines, devoted to experiments in biosciences, condensed matter (magnetic and electronic properties, nanoscience) and materials science among others. Nowadays, two new beamlines are in construction (infrared microspectroscopy and low-energy ultra-high-resolution angular photoemission for complex materials).

In addition to the Controls Section, the Computing and Controls Division has other 3 sections devoted to management systems software, electronics design and support or Information Technology systems and network administration. The mission of the division is to offer services and give support to the beamlines, laboratories and accelerators, as well as to the whole installation to produce high quality experiments. The range of services comprises hardware and software solutions for control systems, personal safety, equipment protection, data acquisition, data analysis, high performance computing, document management, and networks and communications among others.

ORGANIZATION OF THE CONTROL SOFTWARE DEVELOPMENT DURING THE INSTALLATION AND EARLY OPERATION

The Controls Section was setup in the early 2005. During the phases of design, construction and installation the size of the team was progressively increasing until the current 16 members. The section was responsible for the design, installation and commissioning of projects in a wide range of disciplines, for both accelerators and beamlines. The fact of being a single support group for the whole facility showed important competitive advantages during the installation, where peak loads, could be better managed balancing manpower among the different customers. The software design, development, and the protection systems were scheduled in parallel for the accelerators and the beamlines, although the installation and commissioning of the accelerators engendered peak loads, which often took precedence in order to match the milestones and manage the critical paths for the deliveries. The project and the group relied on common tools, technologies, international collaborations and internal-developed transversal libraries and frameworks, crucial for keeping the service support manageable and the service level agreements.

The different products and subsystems were progressively delivered, and run in production. At this stage, the section started to share efforts between the on-going projects and the service support with a number of bugs and defects to fix at the beginning, and later an increasing number of service requests and requests for change. The service desk relies on a ticketing system where every issue is registered with the relevant information to be classified by priority, unit (e.g. an accelerator's group, a specific beamline, etc.), and service (e.g. a subsystem like vacuum, motion, graphical interfaces, etc.). In this scenario, although all products were based on common tools, they were mostly developed by individuals, who ended up taking the responsibility for the related tickets (bugs or new feature requests), and therefore the particular knowledge of the product could not be spread.

Once the Control System has been delivered, during the operation, new features, bug fixes, and advanced functionalities need to be properly scheduled for deployment and commissioning in very restrictive time slots (mostly on shutdowns or "maintenance days"). Furthermore, during operation, the support and maintenance tasks for the running products gain priority

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

USE OF AUTOMATION IN COMMISSIONING PROCESS OF THE UNDULATORS OF THE EUROPEAN X-RAY FREE ELECTRON LASER

Suren Karabekyan, Joachim Pflueger, European XFEL, Hamburg, Germany
Hongxiang Lin, Yongtao Liu, USTC/NSRL, Hefei, Anhui, People's Republic of China
Xuetao Wang, Hisense Co. Ltd., Qingdao, People's Republic of China

Abstract

For operation of the three undulator systems of the European XFEL, a total of 91 undulators are needed and have been produced. For production, magnetic measurements, tuning and commissioning of these devices only two years were foreseen by the project schedule. For these purposes, automated and optimized procedures were needed to accomplish a number of workflows, time-consuming adjustments and commissioning tasks. We created several automation programs which allowed us to reduce the time spent on the commissioning of the control system by an order of magnitude.

OBJECTIVES

All undulators are planned to be commissioned and put into operation in two years, starting from 2013 until the beginning of 2015. Given the fact that European XFEL has three magnetic measurement labs available, 3 weeks were foreseen for commissioning of the control system and the magnetic measurements and tuning of each undulator.

For each undulator, the commissioning procedure includes the following main steps:

1. Installation and adjustment of the linear encoders to achieve an accuracy of ±1 µm.
2. Adjustment of the tilt angles of the undulator girders to better than ±150 µrad.
3. Alignment of the undulator relative to the magnetic measurement bench.
4. Magnetic measurements and field fine tuning.
5. Final measurements and documentation of magnetic properties of undulator.
6. Evaluation and implementation of the correction coefficients for rotary encoders of undulator-motors.
7. Calibration of the temperature sensors.

Items 3 to 5 are part of semi-automatized magnetic measurements and tuning procedures applied to each undulator. To proceed with these tasks, approximately 2.5 weeks are necessary. The items 1, 2, 6, 7, which took without automation 2 - 2.5 weeks, must be automated to be accomplished in less than 3 days.

BASIC APPROACH

The undulator control system is based on industrial components produced by Beckhoff Automation GmbH. Using the TwinCAT software, a programmable logic controller (PLC) can be run either on an industrial PC which is used to control an undulator, or it can be run on a separate Windows PC carrying out the following steps:

- Initialization: Connection to the undulator PLC or to external controllers or devices, loading and initializing of configuration data.
- Data acquisition, depending on the configuration data.
- Data processing and analysis, graphical representation of the analysis results and calculation of required values.
- Formatting the output as required, such as creating ASCII tables, PDF protocols or configuration files for the undulator PLC

ADJUSTMENT OF THE LINEAR ENCODERS

The undulator gap is controlled by four servo motors. Each of them has its own, absolute rotary encoder. This allows an indirect measurement of the gap. A second, more precise system, is directly measuring the gap by means of two absolute linear encoders HEIDENHAIN LC 183 installed on both ends of the undulator girders (see Fig. 1). Both systems can be used alternatively to provide feedback to the motion control system.

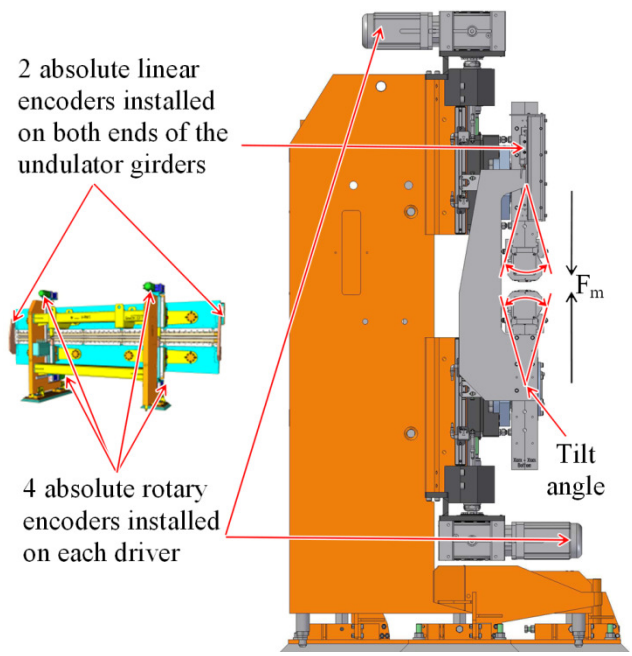


Figure 1: Undulator control and adjustment components.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

INTERFACE MANAGEMENT FOR THE SKA TELESCOPE MANAGER

P.S. Swart*, G.M. le Roux, SKA South Africa NRF, South Africa
A. Marassi, R. Smareglia, INAF-OAT, Italy
S. Vircic, NRC-Herzberg, Canada
S.R. Chaudhuri, TRDDC, India

Abstract

The Square Kilometre Array (SKA) project is currently in the Pre-construction Phase. During this phase, the telescope subsystems are being designed. The Telescope Manager (TM) is a supervisory control and monitoring subsystem in each of the two radio telescopes of the SKA (SKA1-Low and SKA1-Mid). The TM interfaces with a number of diverse telescope subsystems. Interaction between TM and these subsystems is a major source of requirements for the TM. Careful management of TM external interfaces is therefore important. This discussion is a case study of TM interface management. Firstly, how system architectural design aspects like separation of concerns in the control hierarchy reduce telescope complexity with regards to interfaces is discussed. Secondly, the standardisation approach for monitoring and control interfaces to facilitate early elicitation of interface requirements for the TM, and to manage the diversity of interfacing subsystems is discussed. Thirdly, the relations between interface definition and requirements analysis activities, using SysML representations as an example is discussed.

BACKGROUND

The SKA Phase 1 (SKA-1) Observatory will consist of the general headquarters and two radio telescopes: SKA1-Mid in South Africa and SKA1-Low in Australia. Each telescope consists of a number of sub-systems, which are in the SKA project referred to as Elements, of which a TM is one.

There are therefore two TM items, namely SKA1-Mid TM and SKA1-Low TM. For brevity, in this paper the focus is on SKA1-Mid TM (therefore omitting specifics of SKA1-Low TM) because the two telescopes are similar from the TM perspective in its function and interfaces. For less formal reading, "TM" is used when referring to common aspects of these two items, whereas "SKA1-Mid TM" refers to the TM of the SKA1-Mid telescope.

TM OVERVIEW

The TM is a distributed, computer-based supervisory Monitor and Control (M&C) system. TM is responsible for managing [1]:

- astronomical observations,
- telescope hardware and software sub-systems in order to perform those astronomical observations,
- data to support users in achieving operational, maintenance and engineering goals.

For observation management the TM: assists creation, assessment and approval of science proposals; generates schedulable observation units from science proposals; schedules observations; and coordinates execution of the observations to achieve the data products required by the end user.

Telescope management entails configuration and high level control of the telescope Elements and collection of monitoring data which allows the user to monitor the status of Elements, infrastructural support systems and environmental conditions that impact telescope operations.

The Telescope status and configuration data are managed to form a system model that describes the status of the telescope in real-time and historically. To achieve this, the TM collects and stores telescope configuration, telescope dynamic status and environmental data. Stored data are presented to users as current and historic telescope state to assist users during system integration, commissioning, operations and maintenance.

SKA1-MID TM EXTERNAL INTERFACES

This paragraph gives an overview of the external interfaces of the TM, using SKA1-Mid TM as an example (refer [1] par. 1.3). Figure 1 shows SKA1-Mid TM and the systems that it interfaces with.

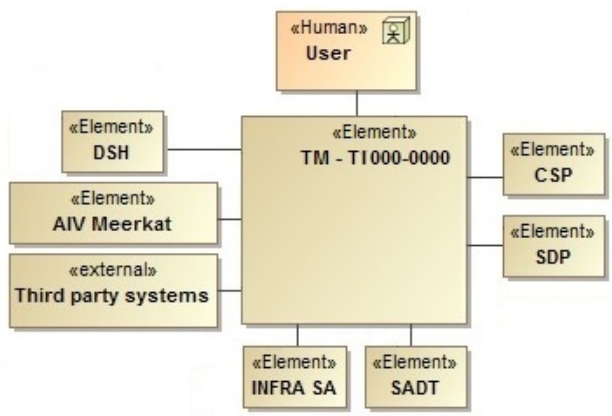


Figure 1: TM relations with SKA environment.

Interfaces with Telescope Elements

The TM has two basic types of interfaces with Elements: data exchange, and physical. The data exchange interface type is categorised into: M&C, general, networking and synchronisation and

*paul@ska.ac.za

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

PYTHON SOFTWARE FOR MEASURING WAVELENGTH AT OPTICALLY PUMPED POLARIZED ION SOURCE (OPPIS)*

P. Kankiya†, J. Jamilkowski, Brookhaven National Laboratory, Upton, NY 11973 USA

Abstract

Often diagnostic tools are packaged with proprietary software and it is challenging to integrate with native environment. The HighFinesse Angstrom Wavemeter used at OPPIS experiment for laser wavelength measurement is controlled using commercial software not supported by RHIC style controls. This paper will describe the integration of such a complex system and use of python for cross platform data acquisition.

INTRODUCTION

The control system environment at BNL's Collider Accelerator Division is a rich repository of control system applications developed in-house and supports complicated supervisory controls and data acquisition tasks, although is limited to supporting these tasks in the Linux environment. Very rarely do we have a requirement to support systems that require software development in a non-UNIX based platforms. We have explored a few solutions in meeting this limitation [1]. In this paper a solution to integrate windows binaries using python is presented. By implementing this solution it will be possible to log data and generate alarms, use synoptic displays etc after gathering measurements from foreign binaries.

SYSTEM IN USE

The polarized beam for RHIC spin physics experimental program is produced in the Optically-Pumped Polarized H-Ion Source (OPPIS), the RHIC polarized H-ion source is upgraded to higher intensity (5-10 mA) and polarization by using a very high brightness fast atomic beam source for enhanced luminosity RHIC operation. Pulsed laser is used to optically pump the rubidium vapour. It is important to measure laser wavelength synchronous to AGS cycle to ensure polarization in desired direction [2]. A high precision device that can measure pulsed laser parameters with high accuracy at a fast rate is needed. To meet this requirement Armstrong meter model shown in Figure 1 was procured with a vision that it will be integrated into RHIC control applications.

Data Acquisition Procedure

Wavelength meter is setup in pulsed operation and externally triggered mode. These control set up tasks are performed by using proprietary software supplied with the hardware. Operator can provide preferred unit of measurement. Operation is started by adjusting the exposure manually or selecting automatic manner.

*Work supported by Brookhaven Science Associates, LLC under Contract No. DE-AC02-98CH10886 with the U.S. DOE

† pkankiya@bnl.gov

For the purpose of integrating measurement results from this device into custom applications, all program values and settings are accessible for user via calls to a dynamic linked library (DLL) called wlmData.dll routines. The wlmData.dll uses shared memory for inter-process communication. The software requires that the server (the Wavelength Meter application) and the clients (any access and control programs) access the same DLL file and not just identical copies which leads to limiting the use of the interface software to one instance at a given time. Elimination of restriction of use of DLL by one client at a time and integrating the whole setup with in-house controls is important to make use of existing infrastructure and correlate data measurements.



Figure 1: High Finesse Wavelength measuring device interfacing a windows computer and a light source.

SOLUTION

There are three common approaches to solving the aforementioned problem:

1. Write your own low level library to communicate with the device.
2. Write an extension in C code to bridge the gap between the libraries and use a cross platform compiler for generating supported executable format.
3. Wrap the library using your **foreign function interface** (FFI) support. Foreign refers to the fact that the functions come from another language and environment.

Because the wlmdata.dll is not an open source file, option one cannot be implemented. Option 2 is more time consuming and limitations has been addressed in another attempt described in paper [1]. Option 3 is considered to be a pragmatic solution. From multiple foreign function

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

CONTROL SYSTEM FOR A DEDICATED ACCELERATOR FOR SACLA WIDE-BAND BEAM LINE

N. Hosoda, T. Fukui, T. Ohshima, T. Sakurai, H. Takebe#, RIKEN/SPring-8, Hyogo, Japan
M. Ishii, JASRI/SPring-8, Hyogo, Japan

Abstract

The X-ray free electron laser facility SACLA has provided opportunities for the exploration of new science. However, these opportunities are severely limited, because only one beam line (beam line 3) is currently available for user experiments. Therefore, there is an urgent need to increase the number of beam lines available. One approach is to install an additional accelerator. We decided to move the SCSS prototype accelerator to an upstream space of beam line 1 in the SACLA undulator hall and reused it as a dedicated accelerator for the beam line. We started an RF conditioning of the dedicated accelerator in October 2014 using MyCC, a MySQL-based temporary data acquisition system compatible with MADOCA. After the RF conditioning, we smoothly transitioned the system from MyCC to MADOCA. The beam commissioning of the dedicated accelerator started in September 2015, and we successfully observed the first extreme ultraviolet free electron laser in October 2015. The control system ensures the coordinated seamless operation of the SACLA accelerator and the dedicated accelerator.

INTRODUCTION

The X-ray free-electron laser (XFEL) facility SACLA (SPring-8 Angstrom Compact X-ray free electron LAser) aims to generate coherent X-ray beams with a wavelength of less than 0.1 nm [1]. It consists of a 400-m-long 8 GeV linear accelerator and a 120-m-long insertion device. In 2011, a commissioning of SACLA was started with beam line 3 (BL3). After the success of generating XFELs, the user experiments at BL3 started in 2012. The SACLA is providing opportunities for the exploration of new science.

Because an XFEL facility employs a linear accelerator, it provides an electron beam for only one beam line at a time. This means that its availability for user experiments is severely limited compared to that of a storage ring-type photon facility such as SPring-8, which has 57 beam lines. Therefore, there is an urgent need to allow a multi-beamline operation of SACLA.

One approach is to deliver an electron beam to multiple beamlines pulse-to-pulse using fast switching magnets [2]. Originally, a DC bending magnet at the end of the accelerator was used to switch an electron beam path between beam lines. We replaced it with a fast kicker magnet and DC twin septum magnets in January 2015. As a result, we successively injected electron beams to beam line 2 (BL2) and BL3 and achieved simultaneous lasing in both beam lines.

Another approach is to install an additional accelerator. Prior to the construction of SACLA, we built a SCSS (SPring-8 Compact SASE Source) prototype accelerator in 2005 to verify the feasibility of lasing using a thermionic gun [3]. We achieved the generation of SASE FEL with a wavelength of 50–60 nm in an extreme ultraviolet (EUV) region in 2007. After the initial purpose was accomplished, it was used for user experiments until 2013. We decided to move the SCSS prototype accelerator to an upstream space of beam line 1 (BL1) in the SACLA undulator hall, and reused it as a dedicated accelerator for BL1 [4].

The SACLA accelerator and the dedicated accelerator share one beam line, BL1. We had to coordinate both accelerator operations, while the independence was maintained as much as possible.

When we designed a control system for the dedicated accelerator, it was natural to use the MADOCA (Message And Database Oriented Control Architecture) control framework, which was already used at the SCSS prototype accelerator and the SACLA [5–7].

The requirements for the control system were as follows. The first priority was not to disturb SACLA user experiments. The RF conditioning was tightly scheduled to start in October 2014 and the beam commissioning was scheduled in September 2015. The human resources available to develop the system were limited. The period in which we could enter the undulator hall to test equipment was also limited. Moreover, the dedicated accelerator was operated at the SACLA control room.

Our strategies were as follows. We reused all control software developed for SACLA. At the start of the RF conditioning, we used a temporary data acquisition (DAQ) system compatible with MADOCA, because some RF components were not installed. After the RF conditioning and a completion of the construction, we switched the DAQ system from MyCC to MADOCA.

In this paper, the status of the dedicated accelerator is reported in terms of its control system.

THE DEDICATED ACCELERATOR FOR BL1

Figure 1 shows a schematic view of the dedicated accelerator for BL1 at the SACLA facility. Although we planned for five beam lines in the undulator hall, only three beam lines had been constructed until now. The bending angle of the transport lines from SACLA to BL1 is 3 degrees and BL1 is located 6 m apart from the SACLA accelerator in parallel. Therefore, a 110-m-long vacant space exists upstream of BL1. The dedicated accelerator was placed in that space.

Present address: Okinawa Institute of Science and Technology (OIST)

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

STATUS OF THE PAL-XFEL CONTROL SYSTEM

Changbum Kim*, Geon-Yeong Mun, Soungyoul Beak, Youngjin Seo, Dong Cheol Shin,
Byoung Ryul Park, Woul-Woo Lee, Jihwa Kim, Heung-Sik Kang, and In Soo Ko,
Pohang Accelerator Laboratory, Pohang, Kyungbuk, Korea

Abstract

Pohang accelerator laboratory (PAL) started an x-ray free electron laser project (PAL-XFEL) in 2011 [1]. In the PAL-XFEL, an electron beam with 200 pC will be generated from a photocathode RF gun and will be accelerated to 10 GeV by using a linear accelerator. The electron beam will pass through undulator section to produce hard x-ray radiation. In 2015, we will finish the installation and will start a commissioning of the PAL-XFEL. In this paper, we introduce the PAL-XFEL and explain present status of it. Details of the control system will be described including a network system, a timing system, hardware control systems and a machine interlock system.

PAL-XFEL

Figure 1 shows a bird eye's view of the PAL facilities which is taken in the April 2015. A circular building in right hand side is the storage ring of the Pohang Light Source (PLS-II) and a new long building on the left hand side is the PAL-XFEL. The building construction was started in 2012 and it was finished at the end of 2014.

The length of the PAL-XFEL building is 1.11km which is consist of a 700 m long linear accelerator, a 250 m long undulator hall, and another 60 m long beam line hall. 10 GeV electron beam can be obtained from the 700 m long linear accelerator and 0.1 nm wavelength hard x-ray can be generated from the electron beam by using the self amplification of the spontaneous emission (SASE).

The total budget is 400 million US dollar and the project period is from 2011 to 2015. Installation of the components of the accelerator was started from the beginning of 2015.



Figure 1: Pohang Accelerator Laboratory.

*chbkim@postech.ac.kr



Figure 2: Pohang Accelerator Laboratory.



Figure 3: Pohang Accelerator Laboratory.

Figures 2 and 3 show the klystron and modulator gallery and inside of the linear accelerator tunnel, respectively. 51 klystrons and modulators and 174 accelerating columns will be installed up to end of 2015.

CONTROL SYSTEM

The PAL-XFEL control covers a large range of hardware and software. An Experimental Physics and Industrial Control System (EPICS) will be used for the control system of the PAL-XFEL. A powerful network system was installed and a large number of Input Output Controllers (IOCs), for example about 700 magnet power supplies, will be connected to a high level control system. An event timing system will be used to send trigger signals to all the equipment and RF distribution system will be ready to supply the 2856 MHz RF signal for S-band RF stations. In addition, a machine interlock system will be prepared to protect important equipment.

Figure 4 shows the network system layout and core switches with Fiber Distribution Frames (FDFs). Duplicated core switches were installed for safety and local switches were connected to them by using 1 Gigabit or 10 Gigabit bandwidth optical fiber.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

XFEL MACHINE PROTECTION SYSTEM (MPS) BASED ON UTCA

S. Karstensen, M. E. Castro Carballo, J. M. Jäger, M. Staack (DESY, Hamburg).

Abstract

For the operation of a machine like the linear accelerator XFEL at DESY Hamburg, a safety system keeping the beam from damaging components is obligatory. This machine protection system (MPS) must detect failures of the RF system, magnets, and other critical components in various sections of the XFEL as well as monitor beam and dark current losses, and react in an appropriate way by limiting average beam power, dumping parts of the macro-pulse, or, in the worst case, shutting down the whole accelerator. It has to consider the influence of various machine modes selected by the timing system. The MPS provides the operators with clear indications of error sources, and offers for every input channel the possibility to set dedicated machine modes to facilitate the operation of the machine. In addition, redundant installation of critical MPS components will help to avoid unnecessary downtime.

This paper summarizes the requirements on the machine protection system and includes plans for its architecture and for needed hardware components. It will show up the clear way of configuring this system - not programming.

Also a look into the financial aspects (manpower / maintenance / integration) is presented.

INTRODUCTION

The European X-Ray Free Electron Laser (XFEL) linear accelerator will bring an electron beam to the energy of up to 20 GeV and use it to generate extremely brilliant pulses of spatially coherent x-rays in an array of undulators using the Self-Amplified Spontaneous Emission (SASE) process. With a design average beam power of 600 kW and beam spot sizes down to few micrometers, the machine will hold a serious damage potential unless countermeasures are taken. To ensure safe operation of the accelerator, dangerous situations have to be detected by closely monitoring beam losses and the status of critical components, and to react appropriately. This is the task of the fast machine protection system (MPS) described in this paper. Several design features of the system have been influenced by experience from existing facilities, particularly the Free Electron Laser in Hamburg (FLASH).

A high flexibility of the MPS is essential to guarantee minimum downtime of the accelerator. In contrast to a storage ring where a beam dump typically implies a time-consuming refill of the machine, a linac offers the possibility to limit the length of the bunch train individually for each macro-pulse. Hence the reaction to failures of subsystems or even parts of the MPS can be much more specific—a dynamic limitation of the total beam power or a selective veto on beam transport into a particular branch of the beamline are possible. Experience with FLASH has also shown that the operation of the machine profits from an MPS whose behaviour can be changed or extended in a simple way. Apart from all of this, overall limitations, like power density, have to be respected.

XFEL ARCHITECTURE

The present chapter provides a brief overview of the XFEL facility (Figure 1). The major tunnel sections accommodate the following systems:

- injector
- linear accelerator (linac)
- beam distribution system
- undulators
- photon beam lines
- experimental stations

These components are distributed along an essentially linear geometry, 3.4 km long, starting on the Deutsches Elektronen-Synchrotron (DESY) laboratory campus in the north west part of the city of Hamburg, and ending in the neighboring federal state of Schleswig-Holstein, south of the city of Schenefeld, where the experimental hall will be located.

In the injector, electron bunches are extracted from a photocathode by a laser beam, focused and accelerated in a radio frequency (RF) gun and a superconducting acceleration module, and directed towards the linac with an exit energy of 120 MeV. After further acceleration, the bunches are longitudinally compressed in two bunch compressors, BC1 and BC2, at energies of 500 MeV and 2 GeV. In the subsequent main linac, the beam is brought to energies of up to 20 GeV (17.5 GeV is the energy foreseen for normal operation of the XFEL) before passing the collimation section. Afterwards, a fast kicker can send single bunches into a beam dump and the beam will be limited. The remaining bunch train can be sent into two undulator lines by the beam distribution kicker with a rise time of less than 20 µs. Each of the undulator lines ends in an electron beam dump, and each of the three main beam dumps is designed to withstand only half of the nominal beam power, i.e. 300 kW.

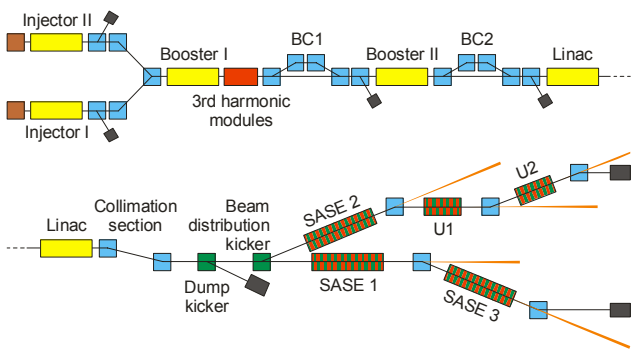


Figure 1: Schematic of the XFEL beam line.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

UPGRADE OF THE BEAM MONITOR SYSTEM FOR HADRON EXPERIMENTAL FACILITY AT J-PARC

Y. Morino*, K. Agari, Y. Sato, A. Toyoda,
KEK, High Energy Accelerator Research Organization, Ibaraki 305-0801, Japan

Abstract

Hadron experimental facility at Japan Proton Accelerator Research Complex (J-PARC) has been improved toward safety operation of high intensity beams. The monitor system of the slow-extraction beam line also has been upgraded to detect abnormal beam extraction. The rate monitor which measured the number of second particles from the production target with the rate of 1 kHz was prepared to detect short pulse extraction. The beam profile monitor was upgraded to read out the data at several times during a beam spill to detect a beam with the narrow width even for a short time. The beam loss monitor was upgraded to monitor beam loss with the rate of ~10 kHz to detect unexpected large loss immediately.

INTRODUCTION

Hadron experimental facility at Japan Proton Accelerator Research Complex (J-PARC) is designed to provide high intensity beams for particle and nuclear physics [1]. A 30-GeV proton beam show-extracted from main ring is injected to a production target (T1) at the hadron experimental facility. The acceleration cycle for the slow-extraction is 6 sec and the spill length is 1.5~2.0 sec.

Figure 1 shows an overview of the hadron experimental facility. In the beam switching yard, there are 33 beam line magnets in 250 m tunnel. In the HD-hall, the proton beams are injected to the production target (T1) and produce secondary particles. The proton beams are transported to the beam dump and absorbed safely. Several kinds of beam monitors are installed in the slow-extraction beam line. The beam monitors involving this paper are shown at Fig. 1.

On May 2013, a proton beam was instantaneously extracted to hadron facility in 5 msec. The short pulse beam molt the production target. After the accident, the beam operation was stopped at the hadron experimental facility. For the recovery of the hadron experimental facility, we upgraded the slow-extraction beam line in many aspects to sustain the abnormal beam extraction. The beam monitor system of the slow-extraction beam line was also upgraded to detect the abnormal beam extraction. The monitor system was developed based on EPICS [2]. The interlock system of the beam line was also improved based on the upgraded monitor system to protect the machines of the accelerator and the beam line.

The rate monitor was upgraded to measure the number of second particles from the production target with the rate of 1 kHz for the detection of short pulse extraction. The readout system of the beam profile monitor was improved to detect a beam with the narrow width even for a short time. The main purpose of the rate monitor and the beam profile monitor is to detect the abnormal beam extraction which may damage the production target. The readout system of the beam loss monitor was improved to detect unexpected large loss immediately. The operation of the slow-extraction beam line was restarted on Apr 2015 successfully.

RATE MONITOR

The rate monitor measures the number of second particles from the production target by using plastic scintillation counters. The rate monitor is installed outside the shielding enclosure of the beam line. The rate monitor was upgraded to read out the particle count with the rate of 1 kHz. Since the number of the second particles is sensitive to the beam intensity, the rate monitor is used to monitor the time structure of the beam intensity. The signals from the scintillators are converted to short pulses with the width of ~60 nsec by a NIM discriminator. The number of the second particles is measured by counting the pulses with a VME module with a FPGA board (GNV-251). The EPICS Input/Output Controller (IOC) runs on Linux on a VME CPU modules (VME-7807).

Figure 2 shows the online monitor display of the rate monitor. The number of the particles per 1 msec is plotted as a function of time. The upper panel shows the example plot in the case of the normal beam extraction. The lower panel shows the predicted plot in the case of the short pulse extraction. In the case of the short pulse extraction, the measured count will increase instantaneously, while the integrated number of the measured count will decrease due to the dead time of the read out system.

The important roles of the rate monitor are the detection of the short pulse extraction and the generation of signals for the interlock system. The upper threshold is set for the number of the particles per 1 msec (instantaneous count). In addition, the upper and lower thresholds are set for the number of the particles per 3 sec (integrated count). In the GNV-251 module, the integrated count is compared with the thresholds for every spill and the instantaneous count is compared with the threshold for every 1 msec. When the particle count crosses the thresholds, the GNV-251 module generates the interlock signal immediately and then the interlock system stops the next beam extraction [3, 4]. The control system of the rate monitor was also developed based on EPICS. The timing of the read out and the threshold values are controlled via EPICS records.

* ymorino@post.kek.jp

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

NONLINEAR SYSTEM IDENTIFICATION OF SUPERCONDUCTING MAGNETS OF RHIC AT BNL*

P. Chitnis#, Stony Brook University, NY, USA
K. A. Brown, Brookhaven National Laboratory, NY, USA

Abstract

The Quench Detection System (QDS) of RHIC detects the Superconducting (SC) magnet quenches by voltage sensing. The real-time voltage across the SC magnet is compared with a predicted voltage from a behavioral model, a deviation from which triggers the quench event and energy extraction. Due to the limitations of the magnet model, many false quench events are generated that affect the RHIC availability. This work is targeted towards remodeling the magnets through nonlinear system identification for the improvement in QDS reliability. The nonlinear electrical behavior of the SC magnets is investigated by statistical data analysis of magnet current and voltage signals. Many data cleaning techniques are employed to reduce the noise in the observed data. Piecewise regression has been used to examine the saturation effects in magnet inductance. The goodness-of-fit of the model is assessed by field testing and comprehensive residual analysis. Finally a new model is suggested for the magnets to be implemented for more accurate results.

INTRODUCTION

The RHIC SuperConducting (SC) magnets store an energy of 70MJ in the form of magnet currents during a full energy run. The SC magnets are susceptible to quenches that lead to the development of tiny resistive zones. An operating current near 5000A (for a dipole magnet) can dissipate this enormous energy at this tiny resistive point causing catastrophic damage.

To safeguard against such failure, Quench Detection System (QDS) is employed. It monitors the SC magnets to detect the developing quenches and sends the magnet power dump signal and beam abort signal to the beam permit system [1]. Voltage sensing is employed for recognizing the developing quenches. The QDS consists of DSPs which store the electrical behavioral model of the SC magnets. The actual magnet output is compared to the model output, and a deviation is sensed as a developing quench, which generates a quench trigger.

The SC magnets exhibit a highly nonlinear behavior due to saturation and hysteresis of steel yoke [2]. Due to mathematically intractable nature of this behavior, the magnet model parameters are manually calibrated, which inhibit the accurate tuning of the model. Also it consumes valuable time when RHIC is running at 4K temperature. Inaccuracies introduce deviation in the model output, which leads to false failures, and resulting in unnecessary machine downtime. Thus to improve QDS reliability, it is necessary that the model truly imitates the SC magnet behavior. The aim of this work is to facilitate automatic generation of accurate magnet models through nonlinear system identification that will improve the reliability and availability of QDS.

Original Magnet Model

The SC magnet circuit's electrical behavior is modeled as a pure inductor with a series resistor. The pure inductor represents the SC magnet and the resistance represents the current leads to the magnet. The model is

$$V_c = L\frac{dI}{dt} + RI$$

Here I is the current through the magnet, L is the nonlinear magnet inductance, R is the lead resistance and V_c is the calculated voltage from the model. This voltage is compared to the observed voltage V_o across the magnet in real time. When a quench develops, additional resistance will appear causing V_o to deviate from V_c. A difference more than 25 mV is triggered as a quench event.

The parameter L is highly nonlinear in nature. It exhibits saturation i.e. its value decreases with increasing current. Also the L *vs.* I curve changes with the change in current ramp waveform. The model stores lookup table for L *vs.* I values, which have to be updated manually every time a new current waveform is introduced. An L *vs.* I lookup table is shown in Fig. 1. This gives us a rough idea of the nonlinear behavior of L.

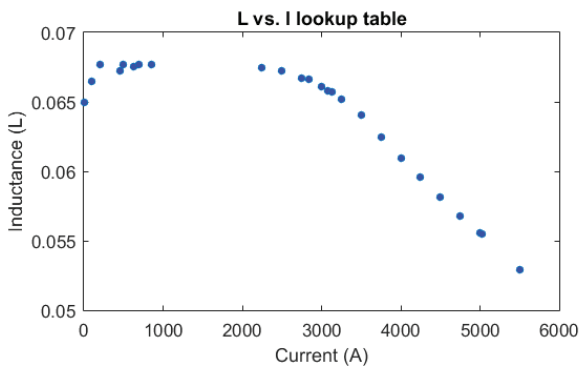


Figure 1: L vs. I lookup table values

*Work supported by Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy.
#prachi.chitnis@stonybrook.edu

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

ALMA RELEASE MANAGEMENT: A PRACTICAL APPROACH

R. Soto, T. Shen, N. Saez, Joint ALMA Observatory, Santiago, Chile
J. Ibsen, European Southern Observatory, Santiago, Chile

Abstract

The ALMA software is a large collection of modules, which implements the functionality needed for the observatory day-to-day operations. The main ALMA software components include: array/antenna control/correlator, submission/processing of science proposals, telescope calibration and data archiving. The implementation of new features and improvements for every software subsystem must be coordinated by considering developers schedule, observatory milestones and testing resources available to verify new software. This paper describes the software delivery process adopted by ALMA since the construction phase and its evolution until these days. It also presents the acceptance procedure implemented by the observatory for validating the software used for science operations. Main roles of the software delivery and acceptance processes are mentioned on this paper by including their responsibility at the different development and testing phases. Finally, some ideas are presented about how the model should change in the near future by considering the operational reality of ALMA Observatory.

OVERVIEW

The software delivery process of ALMA Observatory has passed for several transformations aligned to observatory's project lifecycle. Thus, it changed from a static model (with few and big releases) to a dynamic schema with emphasis at the testing phases and reducing the integration time required for a new release. This paper will present the evolution of the ALMA Software Delivery Process and Release Management by describing the advantages/disadvantages of the models adopted. Final section will give an overview how the model will change in the near future.

SOFTWARE RELEASES IN THE PAST

The ALMA Observatory started its commissioning phase at the Chilean site during the end of 2009. It considered some antennas installed at the high site (5000m) and the deployment of the first quadrant for the baseline correlator. Additionally, the assembly, integration and verification activities (AIV) related to the new array elements delivered by manufacture vendors [1], continued more intensely at the Operation Support Facilities (3000m). This period was very intensive for the computing group, since simultaneous activities had to be supported in parallel. Commissioning process by using direct observing systems (control and correlator software, front-end archive, etc.) was required at the Observatory and, at the other hand, the preparation, integration and testing of pre and post observing software (Phase 1 and Phase 2 proposal submission, time allocation committee support, data processing software, etc.) was demanded as an Early Science task, planned to be started at the end of 2010. In terms of software releases, a cycle of 6 months for the delivery of new features was established, which included capabilities for the observing systems and the proposal handling process as well. These cycles considered several phases (with a timeline predefined) before declare the software as accepted for AIV activities or Early Science observations as described at [2]: *Code Freeze Period, Initial Integrated Testing, Initial Site Testing, Computing Release, Routine Use, Acceptance Testing and Integrated Testing*. The model worked relatively well during early construction stages when the activities were concentrated at the ALMA Test Facilities [3]. Software was commissioned by using prototypes antennas and there was time available with the operational hardware for testing purposes. However, this approach was deficient when commissioning and AIV activities started at the operational site. There were less access to the hardware for testing and more pressure for having new software capabilities working in order to continue progressing into array commissioning.

Preliminary testing was initially executed in a simulated environment, which was not able to reproduce the exact behavior of the operational hardware. So, even a full battery of tests were executed, which considered new functionality and regression tests, there were not enough to deliver a mature software for the site testing. Many bugs were detected using the operational hardware, which increases the cost of correct them [4]. On the other hand, the big amount of features delivered per release also produced additional problems to distill the software. Thus, the stabilization of a new release took about two months after the integration and testing team delivered it. After that, science group should proceed with the routine use and acceptance testing, but the long integration period introduced big delays in the whole commissioning process. Given the current scenery, the integrated computing team looks for changes in the software delivery process adapting the model to the current reality.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

USE INTERRUPT DRIVEN MODE TO REDESIGN AN IOC FOR DIGITAL POWER SUPPLY AT SSC-LINAC

Shi An, Wei Zhang, Kewei GU, Junqi Wu, Xiaojun Liu, IMP, Lanzhou, China

Abstract

SSC-LINAC control system is based on EPICS architecture. The sub control system of digital power supplies is a kind of IOC send and receive custom command via Ethernet and TCP/IP protocol. The old IOC is designed to use period scan mode IOC, and there are so many digital power supplies, that we can't make sure every connect condition of digital power supply is fine. IOC must wait a long time if one of them can't connect correctly and other digital power supply's PV may also be blocked. An IOC that uses interrupt driven mode to avoid the shortcoming was designed. This will be described in this paper.

INTRODUCTION

SSC is a separated-sector cyclotron. To improve the efficiency of HIRFL, a linear accelerator is considered as a new injector for SSC of HIRFL [1]. The SSC-LINAC control system is based on EPICS. There are many sub systems such as power supply control system, vacuum monitor system and so on. The structure of SSC-LINAC control system is shown in Figure 1.

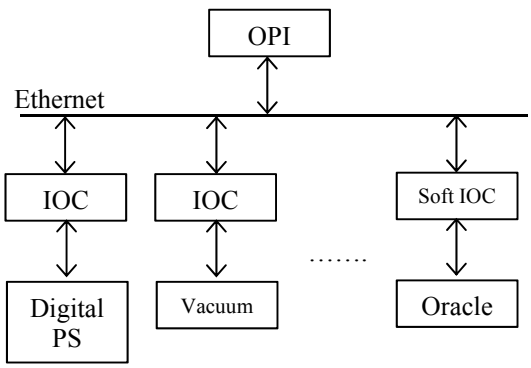


Figure 1: Structure of SSC-LINAC control system.

Almost all of the power supply in the system uses digital power supply. So control interface between IOC and power supply is Ethernet interface. IOC sends and receives predefined string command via the Ethernet with the PS (power supply).

DESCRIPTIONS

This part wills description the old PS IOC design and new design of the IOC, also description implement of the new IOC. Descriptions include two sections and each section is introduced in detail as following. The structure of the PS IOC is shown in Figure 2.

Structure of the Origin PS IOC

The bottom of the structure is digital PS. Above the PS is Ethernet interface and use TCP/IP as low level transfer protocol. In the middle layer is device support of EPICS architecture. The device support has two functions. First function is a periodic that sends a read command string to PS in every one second and receives current state of the PS. The second function is sends a write command string to PS. The top of the structure is record support.

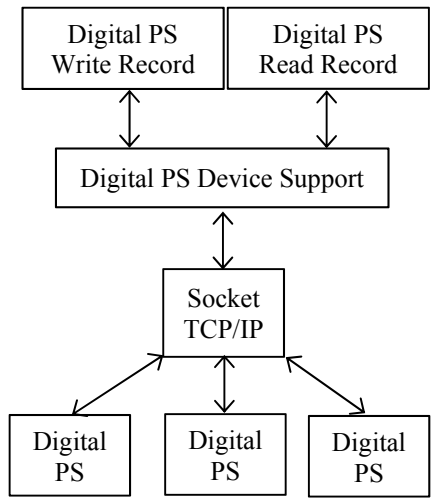


Figure 2: Structure of the PS IOC.

In the real case we create two PVs for every digital PS and there are 78 power supplies in SSC-LINAC. Actually all of PVs are run in one thread in the structure. If just one digital PS is down and all of other digital power supply's PV also be blocked (blocked time depend on the timeout setting). So we must design a new device support to avoid the bad situation.

Design a New Device Support for Digital PS

First switch the Record periodic SCAN ("1 second") to SCAN rate driven from a device specific source ("I/O Intr"). In our case separated thread for every one PS is used.

Second step is writing device support use interrupt model and implement in the following sections.

1. init

```
static long init(int phase)
{
        if(phase==0)
                initHookRegister(&start_ps_thread);
        return 0;
}
```

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

MAGNET CORRECTOR POWER SUPPLY CONTROLLER FOR LCLS-I

S. Babel, K. Luchini, J. Olsen, T. Straumann, E. Williams, C. Yee, B. Lam, SLAC, Menlo Park, California, USA

Abstract

The MCOR 12 is a 16-channel modular architecture, precision magnet driver, capable of providing bipolar output currents in the range from –12A to +12A. A single, unregulated bulk power supply provides the main DC power for the entire crate. Currently the MCORs have a 1000ppm regulation on the B-field. The MCOR controller card upgrades, existing LCLS-I and future LCLS-II needed, controls for Magnet Corrector Power Supplies. The project shifts the existing functionality of the VME based DAC and SAM and an Allen Bradley PLC into a new slot-0 card residing in the MCOR chassis. Elimination of the VME crate and the PLC will free up rack space to be used in future. The new interface card has a long term stability of 100ppm and monitors ground fault currents and various other interlocks for the MCOR power supplies. The controller can interface to EPICS Channel Access and Fast Feedback system at SLAC using two Gigabit Ethernet ports and has an FPGA based EVR for getting "time stamps" from the Event Generator system at SLAC. The EPICS control system along with embedded diagnostic features will allow for enhanced remote control and monitoring of the power supplies.

INTRODUCTION

The MCOR (Magnet Corrector) Series is a multichannel corrector magnet driver system, capable of providing precision bi-polar output currents with minimal zero-crossover distortion. The MCOR design employs a modular architecture, consisting of a rack-mounted crate, with standardized slots for removable power modules, crate controller and diagnostics. A single, unregulated bulk power supply provides the main DC power for the entire crate. The MCOR power module is responsible for converting the unregulated DC bulk power into a precision bi-polar current source suitable for driving corrector magnets and beam line devices. The power modules are controlled using +10 to –10V FS analog command signals sent over the backplane from the control system. There are two types of available MCOR power modules: The MCOR12 (12A FS) and the MCOR30 as shown in Fig. 1 (30A FS). [1, 2] The MCOR30 uses a custom H-bridge configuration, operating in an overlapping volt-seconds feedback mode instead of a straight voltage mode to achieve a higher control bandwidth. The transfer function of input command voltage to output current is determined by resistor values on the programming (PGM) card. Each power module provides two independent measurements of its output current; one that closes the regulation loop, and another that returns an independent monitoring signal to the control system. [1, 2] Both of these independent current measurements are available to the control card for diagnostic purposes. [1, 2] The MCOR12 uses a commercial switch mode servo driver (The AMC 30A8) operating in a voltage feedback mode to regulate the output current. The MCOR 12 system as the name suggests is a 16-channel precision magnet driver, capable of providing bipolar output currents in the range from –12A to +12A. The output current can be adjusted smoothly through zero. [1, 2] The MCOR 12 employs a modular architecture, so that any individual channel is serviceable without disturbing the operation of adjacent channels in the same crate. This feature significantly improves the overall availability of the accelerator, since in most cases the beam lattice can tolerate the loss of a single corrector and continue to operate, but could not handle the loss of an entire crate of correctors during the repair effort. [1, 2]

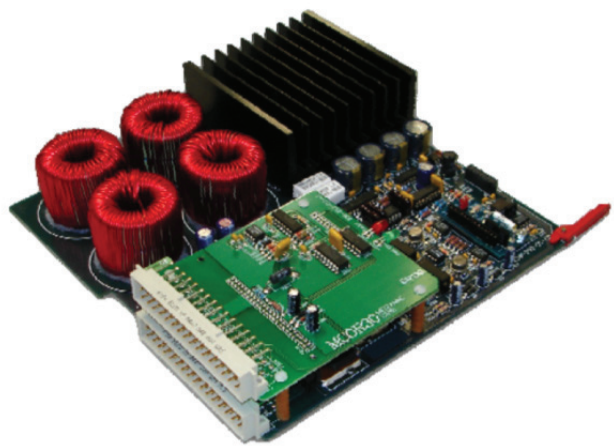


Figure 1: MCOR-30 Power Supply Card.

SYSTEM ARCHITECTURE

The LCLS Fast Feedback Loop Controller wants all MCOR controllers to be controllable at 120Hz (The MCOR Controller card must be able to write set points and read back at 120Hz). The MCOR Controller card must also receive set points in PVs from the Channel Access Network, but it is not necessary that it act on commands from both the Channel Access and the Fast Feedback Network at the same time. LCLS Users want to see the current set point ('desired value'), read back ('actual' value), and status of each channel, regardless of which network the set point command came from. LCLS Users want features like Calibration of the DACs,

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

RENOVATION OF THE CERN CONTROLS CONFIGURATION SERVICE

L. Burdzanowski, C. Roderick CERN, Geneva, Switzerland

Abstract

The Controls Configuration Service (CCS) is a key component in CERN's data driven accelerator Control System. Based around a central database, the service also provides a range of client APIs and user interfaces - enabling configuration of controls for CERN's accelerator complex. The service has existed for 35 years (29 based on Oracle DBMS) [1]. To cater for changing requirements and technology advances there has been substantial evolution of the CCS over time. Inevitably this has led to increases in CCS complexity and an accumulation of technical debt. These two aspects combined have a negative impact on the flexibility and maintainability of the CCS, leading to a potential bottleneck for Control System evolution. This paper describes on-going renovation efforts (started mid-2014) to tackle the aforementioned issues, whilst ensuring overall system stability. In particular, this paper covers architectural changes, the agile development process in place - bringing users close to the development cycle, and the deterministic approach used to treat technical debt. Collectively these efforts are leading towards a successful renovation of a core element of the Control System.

INTRODUCTION

The CERN Control System is a data-driven multi-layer infrastructure including:

- *Low-level hardware and software* – e.g. timing infrastructure, equipment drivers, Front-End Computers (FEC), end-user developed C/C++ binaries representing operational "devices", etc.
- *Middleware layer* – e.g. read/write access to processes running on FECs and Role Based Access Control (RBAC).
- *High-level software* – e.g. high-level settings management, data acquisition and archiving.

The Controls Configuration Service (CCS) helps bind all of the layers together by providing them with complete and coherent configurations that are necessary for the proper functioning of the Control system [2].

The current architecture of the CCS is based on:

- An Oracle database (2-node RAC cluster)
- A set of high-level client Java APIs
- Database level client APIs (PL/SQL interfaces)
- Numerous Graphical User Interfaces based on proprietary Oracle technologies: Application Development Framework (ADF) and Oracle Application Express (APEX).

The database is implemented using a relation model, with approximately 700 domain tables and ~7GB of core domain data (excluding binary, log and history data – which collectively accounts for ~115GB).

The criticality of the service for safe operation of the accelerators chain is high (though not required for their safe shutdown): The CCS is essential for proper accelerator configuration and start-up – especially during Technical Stops when equipment and other components of the Controls System undergo maintenance and upgrades.

The CCS exists for 35 years, during which the scope, architecture, implementation technology and development methodology have kept evolving. In the middle of 2014 the first major service-wide renovation and overhaul has started – marking the beginning of a new chapter in its long history.

RENOVATION STRATEGY

The motivation behind the complete renovation of the service can be summarized as the need to increase service flexibility while lowering total cost of development, and to advance service functionality to the state required by activities planned for the next CERN Long Shutdown (LS2 – scheduled start early 2019).

The cornerstones of the renovation strategy are:

- Suppression of the technical debt accumulated over the years.
- Changes in the overall architecture
- Adaptation of the Lean software development process [3].

All of these aspects are closely related as suppression of technical debt is essential in order to advance the system architecture, while taking proper architectural and design decisions prevent further "erosion" in the system and limit existing technical debt. The adapted software development process facilitates implementing changes: enabling a lower overall cost of development and increased agility. The first two aspects are a mid–to–long–term perspective (from mid–2014 to the start of LS2). The implementation of the Lean software development process is already well advanced and can be considered finished by the end of 2015.

In order to fully understand the context of the renovation efforts it is necessary to briefly look back at the evolution of the CCS scope and technologies used.

Service Evolution

The scope of the CCS was initially limited to the PS (Proton-Synchrotron) complex controls system, meaning that the service and its database were oriented towards a concrete accelerator and its specific control system. The first relational database was introduced in 1986. Over the years the scope grew following the evolution of CERN's accelerator complex [4]. 1995 marks the introduction of graphical user interfaces (GUI) based on Oracle Forms and PL/SQL Web Toolkit (OWA). The first Java based data access API was implemented in 1999 facilitating access for

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

EMBEDDED ENVIRONMENT WITH EPICS SUPPORT FOR CONTROL APPLICATIONS

Y. S. Cheng, Demi Lee, C. Y. Liao, C. H. Huang, K. T. Hsu
National Synchrotron Radiation Research Center, Hsinchu 30076, Taiwan

Abstract

System on a chip (SoC) is widely used in embedded environment. Current generation SoC commercial products with small footprint and low-cost have powerful in CPU performance and rich interface solution to support many control applications. To deal with some embedded control applications, the "Banana Pi" which is a card-size single-board computer and runs Linux-based operation system has been adopted as the EPICS IOC to implement several applications. The efforts for implementing are summarized in this paper.

INTRODUCTION

Using a circuit board to implement functions as a computer is called SBC (single-board computer) [1]. Its applications cover in telecommunications, industrial control, blade and high density servers, and lately laptops and mini-PCs, etc. Thanks to latest generation SoC technology, putting all major functionality into an integrated chip, educational used credit-card size SBC [2] likes the Raspberry Pi (RPi) and BeagleBone Black (BBB) are highly successful products. The Banana Pi (BPi) is the latest product of such category with powerful CPU, low power consumption SBC indeed, and the area of circuit board is only as credit card size.

The Banana Pi which design idea is similar to the RPi-style SBC, and it is a fork of the RPi project using different components while maintain compatibility as much as possible. Moreover the Banana Pi is added the functions of SATA interface, infrared transmission, microphone, USB-OTG ports, power button, reset button, etc. Then the BPi has 26-pin/40-pin GPIO which is compatible with the RPi. The A20/31 SoC as CPU/GPU, 1GB DDR3 memory and Gigabit Ethernet connection are applied on the Banana Pi. The hardware specification of Banana Pi is shown as Table 1 [3-4]. Linux-based OS can be worked well on the Banana Pi.

Table 1: Hardware specification of the Banana Pi

	Banana Pi M1/M1+/M2
CPU	A20 ARM Coretx-A7 1GHz Dual-Core A31S ARM Coretx-A7 1GHz Quad-Core
Memory	1GB DDR3 DRAM
Network	1Gbps Ethernet RJ45, Wi-Fi
Storage	SD card slot (up to 64GB), Extensible with SATA interface
I/O	GPIO, UART, I2C bus, SPI bus with two chip selects, CAN bus, ADC, PWM, +3.3V, +5V, GND
OS	Debian, Bananian, Lubuntu, Android ...

The EPICS (Experimental Physics and Industrial Control System) [5] is a set of open source software tools, libraries and applications developed collaboratively and used to create distributed soft real-time control systems for scientific instruments such as particle accelerators. Many facilities have good practical experiences for the EPICS and adopt it as accelerator control systems. Many resources and supports are available as well as numerous applications for accelerator have been developed.

The TPS (Taiwan Photon Source) control system of 3 GeV synchrotron light source is also based on the EPICS framework [6]. The EPICS toolkit provides standard tools for display creation, archiving, alarm handling, etc. The big success of EPICS is based on the definition of a standard IOC (Input Output Controller) structure together with an extensive library of driver software for a wide range of I/O cards. The EPICS framework which has various functionalities is employed to monitor and to control on embedded applications of accelerator system.

BANANA PI AS EPICS IOC

The stability and performance of Banana Pi is enough as the EPICS IOCs for specific control applications. The EPICS framework can be built on the Linux-based Banana Pi successfully [7]. At the TPS project, some control functions, such as frequency divider, direct digital synthesizer, radiation-sensing reader, alarm announcer, etc., are implemented by use of the Banana Pi platforms with EPICS support. The efforts for implementing are summarized as followings.

Software Architecture

To implement the Banana Pi as the EPICS IOC for specific control applications, the EPICS base and modules are necessary to be set up on the Banana Pi platform which operation system can be the Debian or Ubuntu Linux. The device driver of SPI (Serial Peripheral Interface) bus is built for communicating with DAC/ADC modules, and the device support interface is also developed as the glue between the EPICS records and device drivers. The EPICS records support with databases are created according to the specific functions. The application module, such as “autosave” function, is installed for logging setting parameters values and recovering last setting parameters values automatically when the IOC is start-up. Based on the EPICS PV (Process Variable) channel access, the archive server is set up to record various parameters variations for long time observation, and the PHP webpage can be developed to show the status information. At the client console side, the operation interfaces are created by used of the EDM, CS-Studio, etc. to control and monitor via EPICS PV channel

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

LLRF CONTROLS UPGRADE FOR THE LCLS XTCAV PROJECT AT SLAC*

S.Condamoor#, Y.Ding, P.Krejcik, H.Loos, T.J.Maxwell, J.J.Olsen
SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

Abstract

SLAC's Low Level Radio Frequency (LLRF) controls software for the S-Band deflecting structures needed to be upgraded significantly when a new X-Band transverse deflecting cavity (XTCAV) was installed downstream of the LCLS undulators in Spring 2013 to assist in FEL diagnostics such as characterizing the temporal profile of X-ray pulses that vary shot-to-shot. The unique location of the XTCAV in the beamline posed several challenges. A new design of the Modulator and Klystron control Support Unit (MKSU-II) for interlocking was added at the XTCAV controls station that required new software development. The timing setup was also different from the rest of the Linac. This paper outlines the LLRF controls layout for the XTCAV and discusses the manner in which the challenges were addressed. XTCAV has now become a successful tool for gathering data that enables reconstruction of X-ray FEL power profiles with greater resolution.

XTCAV FOR DIAGNOSTICS

XTCAV is a powerful diagnostic tool for SLAC's Linac Coherent Light Source (LCLS) Free Electron X-ray Laser (FEL). This is a successor to SLAC's historic S-Band TCAV[1], two of which are still used in LCLS for electron beam bunch length measurements. XTCAV provides useful FEL diagnostics without interfering with LCLS operations since it is installed downstream of the LCLS undulators that are the source of the x-ray pulses. XTCAV deflects spent e-beam bunches after they have contributed to the lasing process just after they exit the undulators. Energy spectrometer is then used to capture and analyse in detail the profile of this spent e-beam stream on a bunch-by-bunch basis. Comparison of the longitudinal profile of these electron bunches when they are lasing against when lasing is suppressed allows us to reconstruct not only the longitudinal profile of the incoming electron beam but the time-resolved profile of the X-ray pulses generated during the lasing process. Figure 1 shows the general layout of this diagnostics system.

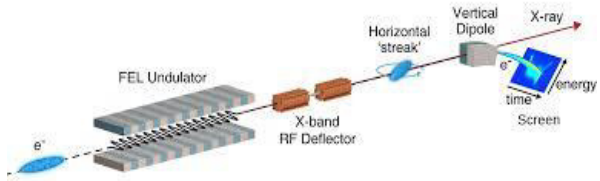


Figure 1: The system includes the transverse deflector "XTCAV", the magnetic spectrometer and the Ce:YAG screen located downstream of the FEL undulators. Horizontal streaking is followed by a vertical-bend dipole magnet for measuring the energy spectrum. A camera captures the transverse images of the electron beam density distribution on the diagnostic screen. (Image: Y.Ding et al., PRSTAB 14, 2011).

XTCAV'S HIGH POWER RF

The SLAC-designed X-Band RF cavity shown in Figure 2 below is powered by SLAC's XL4 50 MW X-Band Klystron that deflects the beam by giving it a 48 MV kick[2]. The Klystron operates at 120 Hz and delivers pulses that can be up to 1.5 uS long. The structure fill time is 200 nS. A legacy 6575 modulator was reused from SLAC, the controls for which were upgraded recently.

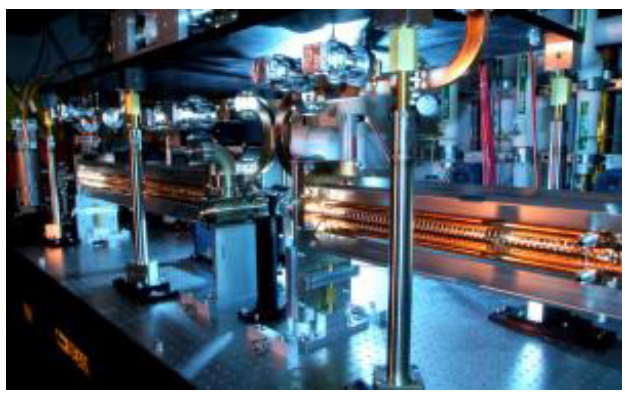


Figure 2: Two one-meter long X-Band RF structures are in the beamline's path towards the e-beam dump. (Image: P.Krejcik/SLAC).

XTCAV CONTROLS

The Controls hardware for XTCAV is located in a service building above the cavity in the tunnel. In addition to the klystron, this building houses the modulator, LLRF controls hardware, power supplies, waveguides, vacuum, temperature and water controls. The LLRF phase reference is derived from the 476 MHz Linac Reference via a heliax cable from the tunnel. Timing and trigger generation for controls is obtained with an Event Receiver

* Work is supported by the U.S. Department of Energy, Office of Science under Contract DE-AC02-76SF00515.
#scondam@slac.stanford.edu

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

FAST WIRE SCANNER UPGRADE FOR LCLS

J. M. D'Ewart#, M. Campell, P. Krejcik, H. Loos, K. Luchini, SLAC, Menlo Park, CA 94025, USA

Abstract

Wire scanners are a main diagnostic tool for transverse beam size and emittance measurements at LCLS. The original SLAC wire scanners were not optimized for speed (taking minutes to scan), and can't perform at the desired level of position resolution necessary for measuring LCLS' small beam size. A new fast wire scanner, based on a dc linear servo motor, has been designed and installed in the LCLS. The new fast wire scanner has several advantages over the original wire scanner: scan times are reduced from minutes to seconds while minimizing wire vibrations. Rather than counting open-loop step pulses, the new fast wire scanner uses real time position capture for beam synchronous sampling of the wire position, enhancing beam profile accuracy.

INTRODUCTION

The primary purpose of the wire scanner system is to provide a high resolution measurement of electron beam profile averaged over many shots. Profile measurements can also be used to determine transverse beam emittance.

The LCLS wire scanner system had been based on a drive system consisting of a stepper motor and a ball screw [1]. A 10X gear reducer was added to the system reducing wire vibration, but also reduced maximum wire speed [2]. A linear variable differential transformer (LVDT) was used to calibrate the drive system. The motion was operated in open-loop with no feedback about actual drive or wire position.

The design goals of the fast wire scanner are: reduce scan time from minutes to seconds, minimize vibration during the scan, improve position resolution of the wire scanner motion to 1μm or better, incorporate real-time readback of the wire position through the data acquisition system, and to integrate the wire scanner into the LCLS EPICS control system [1].

The advent of the LCLS-II project at SLAC has also resulted in the need for a wire scanner that can move at speeds sufficient to prevent wire destruction during CW beam operation.

For the nominal LCLS-II beam parameters of 100 pC charge per bunch, a transverse emittance of 0.5 μm, and a bunch repetition rate of 0.6 MHz the minimum wire speed to avoid damage is calculated to be 0.34 ms^{-1}.

FAST WIRE SCANNER

The fast wire scanner, shown in Fig. 1, uses an external actuator for linear motion. The drive system is based on the LinMot dc linear motor [3]. The wires are mounted on an interchangeable card that holds an x, y and u wire at 45° to the beam. A scan of all 3 wires can be done in a single motion of the scanner before returning to the home position. An increment linear encoder is attached to the drive system and limit switches are used at the end of travel. The linear motor eliminates the discrete steps of a stepper motor and the linear encoder allows for closed-loop servo control.

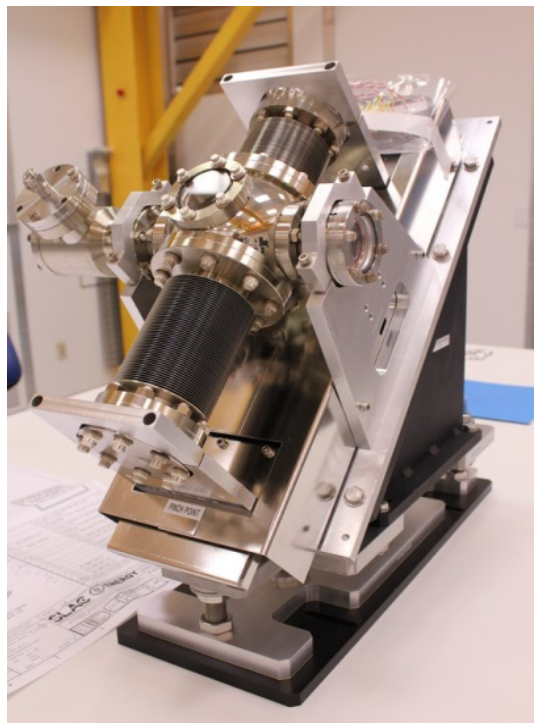


Figure 1: Wire scanner mounted at 45°.

CONTROL SYSTEM

Fast Wire Scanner Control System Architecture

The fast wire scanner control system architecture is based on the original LCLS VME wire scanner architecture. The original LCLS wire scanners use a VME crate with a MVME 6100, MRF EVR, CAEN V965 QDC board, Isig high voltage power supply module, and Hytec IP8601 stepper motor controller. The fast wire scanner uses an additional Prodex MAXv VME motion controller for servo motion control. A MAXv breakout chassis and LinMot panel are also installed into the control rack to support the fast wire scanner. The MAXv breakout chassis separates each of the 8 axis control and I/O signals and interfaces to the LinMot panel. The LinMot panel houses the servo drive as well as drive and logic power supplies. This architecture allows a rolling upgrade for current LCLS wire scanners. A block

*Work supported by the U.S. Department of Energy under contract number DE-AC02-76SF00515
#mdewart@SLAC.Stanford.EDU

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

IMPROVING THE COMPACT MUON SOLENOID ELECTROMAGNETIC CALORIMETER CONTROL AND SAFETY SYSTEMS FOR THE LARGE HADRON COLLIDER RUN 2

D. Di Calafiori, G. Dissertori, L. Djambazov, O. Holme, W. Lustermann, ETH Zurich, Switzerland
P. Adzic, P. Cirkovic, D. Jovanovic, University of Belgrade, Serbia
S. Zelepoukine, University of Wisconsin-Madison, USA

Abstract

The first long shutdown of the Large Hadron Collider (LS1, 2013-2015) provided an opportunity for significant upgrades of the detector control and safety systems of the CMS Electromagnetic Calorimeter. A thorough evaluation was undertaken, building upon experience acquired during several years of detector operations. Substantial improvements were made to the monitoring systems in order to extend readout ranges and provide improved monitoring precision and data reliability. Additional remotely controlled hardware devices and automatic software routines were implemented to optimize the detector recovery time in the case of failures. The safety system was prepared in order to guarantee full support for both commercial off-the-shelf and custom hardware components throughout the next accelerator running period. The software applications were modified to operate on redundant host servers, to fulfil new requirements of the experiment. User interface extensions were also added to provide a more complete overview of the control system. This paper summarises the motivation, implementation and validation of the major improvements made to the hardware and software components during the LS1 and the early data-taking period of LHC Run 2.

INTRODUCTION

The Compact Muon Solenoid (CMS) [1] is a general-purpose particle physics detector built for the Large Hadron Collider (LHC) [2] at CERN. The CMS detector is composed of several types of sub-detectors: trackers, calorimeters and muon chambers. Based on the technical specifications and requirements of each sub-detector, custom Detector Control Systems (DCS) were designed, implemented and commissioned. A central CMS DCS [3] infrastructure provides general and common services, as well as the unification of controls and monitoring of all sub-detectors. A central CMS Detector Safety System (DSS) [1] is responsible for all off-detector safety-related matters. In addition, individual sub-detector safety systems are also available to provide extensions to the CMS DSS to guarantee the on-detector safety and integrity. The experience acquired during the first LHC long run (March/2010 – February/2013) ensured a better understanding of the detector operation and consequently the identification of possible fields for improvements to the DCS/DSS software and hardware components. The first long shutdown (February/2013 – June/2015) of the accelerator provided an opportunity for the implementation of most of these improvements, systems extensions and the preparation of even more robust, efficient and reliable systems for the LHC Run 2. This paper discusses the main improvements made to both DCS and DSS of the CMS Electromagnetic Calorimeter (ECAL) [1] throughout the periods described above. Focus is given to hardware upgrades and system consolidation.

Full details of the architecture of the CMS ECAL DCS have been reported previously [4,5].

CMS ECAL DCS OVERVIEW

This section presents a general description of the CMS ECAL DCS software and hardware. A simplified block diagram of the complete system is illustrated in Fig. 1.

Software

The CMS ECAL DCS application has been developed with a Supervisory Control and Data Acquisition (SCADA) software package named SIMATIC WinCC OA from SIEMENS, formerly Process visualization and control system II (PVSS II) from ETM. To facilitate the development process, the JCOP Framework [6], developed by CERN to provide extensions to WinCC OA for the high-energy physics domain, is extensively used. Through a Finite State Machine (FSM), manual and automatic commands are issued to lower nodes and node states are propagated and summarised upwards. Furthermore, the main control and state monitoring of all sub-detectors are accessible to the CMS DCS via the FSM to enable global CMS control.

The partitioning of the control nodes follows the sub-detectors architecture, in this case: CMS ECAL Barrel (EB), CMS ECAL Endcaps (EE) and CMS ECAL Preshower (ES).

Hardware

The DCS hardware comprises all devices used for connection to the low voltage and bias voltage powering systems, the electronics for environment monitoring and for additional support services, such as remote crate power cycle units and bias voltage current measurements.

Due to software changes at the CMS DCS level, presented in the section "*Controls Software Redundancy*", all field buses are being converted to Ethernet, allowing their physical decoupling from the DCS servers.

Spare parts and, in some cases, complete spare units are available to provide full and optimal support in case of hardware failures.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

EXPERIENCES AND LESSONS LEARNED IN TRANSITIONING BEAMLINE FRONT-ENDS FROM VMEBUS TO MODULAR DISTRIBUTED I/O

I. J. Gillingham, T. Friedrich, S. C. Lay, R. Mercado,
Diamond Light Source, Oxfordshire, UK

Abstract

Historically Diamond's photon front-ends have adopted control systems based on the VMEbus platform.

With increasing pressure towards improved system versatility, space constraints and the issues of long term support for the VME platform, a programme of migration to distributed remote I/O control systems was undertaken.

This paper reports on the design strategies, benefits and issues addressed since the new design has been operational.

INTRODUCTION

Diamond Light Source is a 3 GeV third-generation light source with a 561m storage ring (SR), a full-energy booster (BR) and a 100 MeV pre-injector Linac [1]. The photon output is optimised for high brightness from undulators and high flux from multi-pole wigglers. The current operational state includes 28 photon beamlines, with a further two beamlines in an advanced stages of design and construction.

In 2010, Diamond adopted a new control system design for future beamlines, adopting fieldbus distributed control as a viable alternative to the long standing VMEbus architecture [2]. In planning for the present phase of photon beamlines, in 2015, the control system architecture and choices of hardware used was revisited.

PREVIOUS CONTROL SYSTEM ARCHITECTURES

The original accelerator and beamline control systems at Diamond are based on VME systems. To support the interface requirements of the equipment, a range of I/O modules based on Industrial Pack (IP) modules (ADC, DAC, Serial, DIO) and VME modules (IP carrier, motion, scalar and timing) were used. The field signals are interfaced via either rear transition modules or front-panel connections. A VME microprocessor (MVME5500) runs VxWorks and EPICS to serve up the control information to client applications.

REASON FOR CHANGE

The original control system architecture served well for the existing accelerators and beamlines; however it was defined a number of years ago, so in the context of the recent and future phase of beamlines the opportunity to reconsider the standards was taken. In doing so, it is clear that not all the hardware capability of VME was required for beamline control; neither was the use of a hard real-time operating system such as VxWorks. It was also apparent that most I/O functionality required for control of beamline and accelerator equipment can now be realised through Ethernet-attached I/O. There is also now good infrastructure for developing and managing Linux based EPICS IOCs on a PC architecture.

OUTLINE REQUIREMENTS FOR PHOTON BEAMLINES

In considering the requirements for photon beamlines, the following technical systems were identified:

- Motion control
- Vacuum instrumentation and other serial devices
- Video cameras
- Analogue and digital signals
- Programmable logic controllers
- Timing signals

The disturbance to beamline operation through restarting an IOC, should be minimised.

The I/O associated with the control system should be located close to the equipment being interfaced; i.e. for signals located in experimental and optics hutches, the I/O modules should be co-located in these areas. However, this is constrained by the possibility of radiation-induced damage to I/O in the optics hutches of high energy (~100keV) beamlines or Storage Ring (SR) vault and by the space available in the some beamline hutches.

NEW CONTROL SYSTEM ARCHITECTURE

Each EPICS IOC runs on a 1U Linux PC located within the relevant Controls Instrumentation Area (CIA). The Linux servers have several physically separate network connections to support the different systems.

Linux Server

The IOCs are standard Dell 1U machines (currently R320), without a hard-disk installed, that boot from an image over the network into a RAM based disk.

The current operating system is 64 bit RedHat Enterprise Linux 6, which has real-time extensions compiled in to facilitate determinacy for the EtherCAT driver, when used at high bandwidth.

On boot, the server picks up its host-name and IP address via DHCP, determined by the MAC address of eth0. The DHCP server also points to the TFTP server which provides the kernel and ramdisk image used to boot the machine.

A startup-script in the host's soft IOC directory is used to configure the server. For instance, network interface

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

DETECTOR AND RUN CONTROL SYSTEMS FOR THE NA62 FIXED-TARGET EXPERIMENT AT CERN

Piotr Golonka, CERN, Geneva, Switzerland
Riccardo Fantechi, CERN, Geneva, Switzerland
Manuel Gonzalez-Berges, CERN, Geneva, Switzerland
Fernando Varela, CERN, Geneva, Switzerland
Valeri Falaleev, JINR, Dubna, Russia
Nicolas Lurkin, University of Birmingham, Birmingham, UK
Ryan Frank Page, University of Bristol, Bristol, UK

Abstract

The Detector and Run Control systems for the NA62 experiment, which started physics data-taking in the autumn of 2014, were designed, developed and deployed in collaboration between the Physics and Engineering Departments at CERN. Based on the commonly used control frameworks, they were developed with scarce manpower while meeting the challenge of extreme agility, evolving requirements, as well as integration of new types of hardware. The paper presents, for the first time, the architecture of these systems and discusses the challenges and experience in developing and maintaining them during the first months of operation.

INTRODUCTION

NA62 is a fixed-target experiment at CERN aiming at high-precision measurements of rare Kaon decays, currently taking physics data in its second run. Operation of the detector and its infrastructure is supervised with dedicated control systems, the most complex being the Detector Control System (DCS) and Run Control (RC). The two are similar in many aspects and share common infrastructure, as well as development and maintenance workflows. The Run Control has already been introduced in [1,2]. This paper focuses on the - often neglected - topic of complete architecture, including infrastructure, project organization and development life cycle.

The NA62 detector is composed of a set of sub-detectors developed and maintained by semi-independent groups. Some equipment is reused from previous experiments (such as the Liquid Krypton Calorimeter), while others must be built anew. Certain sub-detectors or their parts were not ready for integration before the first run, and they are added progressively In consequence, the development, commissioning and maintenance of control systems is a *process* that spans over years. This is unlike many turn-key control systems in industry. Assuring consistency and long term maintenance in the project, taking into account a high-turnover of people in experimental collaboration is one of the primary challenges.

REQUIREMENTS

The DCS and RC systems for NA62 resemble those of other CERN experiments. Relative to the LHC experiments they are an order of magnitude lower in number of devices, yet their complexity remains similar.

The systems have to allow for operation and monitoring of a variety of devices, react on value changes, evaluate alarm conditions and notify the operators about the alarms. The operators need to be presented with an alarm screen, as well as synoptic views that summarize the state of parts and subsystems of the experiment, and be able to drill down and quickly identify or locate the required device. They need to easily command individual devices, as well as groups of them (e.g. complete subdetectors)

The experts need to tune the set-points and alarm threshold, and bring up the trend plots of historical values. Reconfiguration should be applied to a selected device or using configuration sets (*recipes*) reflecting the run modes of the experiment. The experts must be able to detect and diagnose malfunctions and communication problems, and dynamically reconfigure the system, e.g. re-wire faulty sensors or mask alarms.

The history of acquired data (*conditions*) needs to be recorded to an Oracle database to be used in physics analysis. CERN computer security rules need to be applied. Robust infrastructure and hardware should be employed to guarantee round-the-clock operation during 8 months of data taking with minimal human supervision.

OPERATION

The independent operation of the DCS and RC is performed on dedicated consoles in the NA62 Control Room. The User Interface window requires authentication with CERN credentials, and enables access to shifter or expert operations, on various parts of detector, depending on privileges defined by administrators.

Standard shifter operation employs two tools: the alarm screen, displaying a time-sorted list of anomalous incidents ("alarms") and the hierarchical control supervision screen. The latter, also known as "FSM" (Finite State machine), abstracts out the hardware to present a tree view of the detector's parts and subsystems. Hardware states are evaluated and represented by colour-codes and text labels ("*ERROR*", "*RAMPING*"). These states are propagated in the upstream direction of the tree using summarization logic. This allows for rapid root-cause identification by expanding the coloured tree view. High level commands ("SWITCH_ON", "CONFIGURE") are broadcast down the tree allowing control of large parts

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

DATABASE ARCHIVING SYSTEM FOR SUPERVISION SYSTEMS AT CERN: A SUCCESSFUL UPGRADE STORY

P. Golonka, M. Gonzalez-Berges, J.Hofer, A. Voitier
CERN, Geneva, Switzerland

Abstract

Almost 200 controls applications, in domains like LHC magnet protection, cryogenics and vacuum systems, cooling-and-ventilation or electrical network supervision, have been developed and are currently maintained by the CERN Industrial Controls Group in close collaboration with several equipment groups. The supervision layer of these systems is based on the same technologies as 400 other systems running in the LHC Experiments (e.g. WinCC Open Architecture, Oracle). During the last two-year LHC Long Shutdown 1, the 200 systems have been successfully migrated from a file-based archiver to a centralized infrastructure based on Oracle databases. This migration has homogenized the archiving chain for all CERN systems, and at the same time has presented a number of additional challenges. The paper presents the design, the necessary optimizations and the migration process that allowed us to meet unprecedented data-archiving rates (unachievable for the previously used system), and liaise with the existing long-term storage system (LHC LoggingDB) to assure data-continuity.

INTRODUCTION

Storage and retrieval of historical process values and alarms is one of the invisible yet essential tasks of SCADA systems, and is always delegated to a dedicated *historian* or *archiver* tool. WinCC OA [1] offers three technologies for archiving: the file-based archiver ("valarch"), the simple database logging facility ("DBLogger", available as of version 3.13) and the high-performance, scalable Oracle database archiver ("RDB Archiver").

Even though a number of large control systems (Detector Control Systems for the LHC experiments) made use of the RDB archiver since the beginning, the control systems for infrastructure and LHC services initially chose to use the file archiver. The divergence in the choice of archiving technology stems from the requirements. The detector control system had to have the historical values available as the so called "conditions" in an Oracle database to make the data analysis possible, and it was this requirement that drove the collaboration between CERN and ETM to develop and optimize the database archiver. On the other hand, the priority for the LHC and infrastructure systems was the use of a stable, proven technology to maximize the reliable operation of the systems.

With an increasing number of control systems to deliver and maintain, and the need to scale-up applications to meet the challenges of Run II of the LHC, it became clear that use of the file-based archiver reached its limit, and a migration to the more performant and centralized archiving architecture became a necessity. In addition, the centralized management of archived data offers several advantages such as better tools for data analysis, simplified management, easier handling of the SCADA server nodes, etc.

In the paper we present the process of migrating the 200 applications of the SCADA Application Service from file-based archiver to Oracle database archiver: the initial requirements, thorough testing and optimisation phase and proof of scalability, redesign of the long-term data storage architecture, actual migration effort. We will then summarize the current state of this enterprise.

EVOLUTION OF ARCHITECTURE

To enable the analysis of data from the control systems for the LHC, data needs to be transferred to a long-term storage Oracle database called the LHC LoggingDB designed to store large amount of data throughout the life time of the LHC [2].

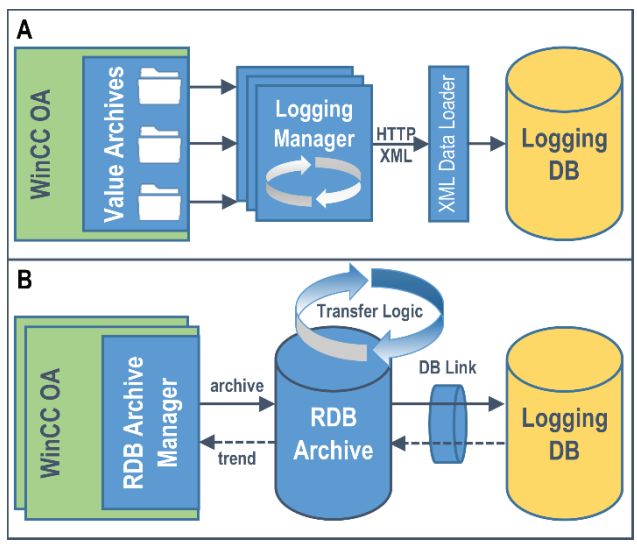


Figure 1: Archiving architecture: (A) with old logging manager; (B) with Oracle archiver and DB-link data transfer.

The evolution of the architecture for archiving is presented in Figure 1. System and process data retrieved by WinCC OA systems used to be fed to the LoggingDB using a dedicated program (Logging Manager) and complex infrastructure (XML Data Loader). This transfer was repeated every couple of minutes, and hence the data was available in the LoggingDB with a lag of

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

SIS18 UPGRADE: THE FAIR COMPLIANT RENOVATION OF THE DATA ACQUISITION SYSTEM FOR PARTICLE DETECTORS

R. Haseitl, H. Bräuning, T. Hoffmann, K. Lang, T. Milosic
GSI Helmholtzzentrum für Schwerionenforschung GmbH, Darmstadt, Germany

Abstract

In preparation of FAIR, several well-established beam instrumentation systems of the GSI heavy-ion synchrotron SIS18 and its connected high-energy beam transfer lines (HEBT) have to be modernised. In this contribution, the data acquisition upgrade of particle detectors such as ion chambers and plastic scintillators is described. This covers the replacement of outdated custom-built readout- and control hardware by modern FMC (FPGA mezzanine card) based I/O hardware, new multi-channel high voltage power supplies and a new data acquisition system (DAQ) for the VME based scalers. The latter will replace the old Kylix-based ABLASS software by LASSIE (Large Analog Signal and Scaling Information Environment) to fit into the new FAIR control system concept. LASSIE is based on FESA (Front End Software Architecture). FESA was originally developed by CERN and enhanced by GSI-specific modifications. Furthermore, the new particle detector DAQ will be able to take full advantage of the new FAIR timing system which is based on the White Rabbit protocol.

INTRODUCTION

For the upcoming FAIR project, many parts of the existing GSI facility have to be modernised. The new control system for FAIR is already operational and currently under test by different GSI departments. It is based on the FESA framework [1] which originally was created by CERN and is now being developed in a collaborative effort with GSI. FESA controlled devices are automatically integrated in the control system and can be accessed in a standardised way. If existing hardware has to be replaced or new hardware is installed, a FESA class is programmed to provide future-proof operation. Recently, a renovation project was started to replace outdated high voltage power supplies and to upgrade the readout software called ABLASS [2] of the SIS18 and High Energy Beam Transfer (HEBT) particle detectors.

HIGH VOLTAGE CONTROL FOR PARTICLE DETECTORS

Several particle detector systems need high voltages (HV) to operate. Currently three main types of particle detectors are used by the GSI beam instrumentation (BI) group in the synchrotron and the HEBT: secondary electron monitors (SEM), ionisation chambers (IC) and plastic scintillators (SC). It is possible to mount these three detectors onto a single pneumatic actuator which is called Particle Detector Combination (PDC). Along the existing synchrotron and the HEBT 13 PDCs are installed. Besides the PDCs, also other devices like beam loss monitors or profile grids need HV resulting in over 120 required channels. The different detectors need specific voltages ranging from -1500V to +1500V.

The existing HV system is based on SY127 40-channel crates by the company CAEN [3]. The crates have been in operation for over 20 years and replacement modules are slowly phasing out. Moreover they are controlled via the slow CAENET bus which can delay new settings for HV channels.

Succeeding HV systems from the companies Wiener [4], iseg [5] and CAEN are already installed at GSI and prepared for operation. These systems provide Ethernet enabled controllers and various HV modules with different voltage and current ranges.

Although both systems are controlled over Ethernet, they use different protocols: The crates from Wiener, called MPOD, are accessible via SNMP (Simple Network Management Protocol). The HV modules for MPOD are manufactured by the company iseg and are also fully controllable via SNMP. At GSI, two different MPOD systems are used: The 'MPOD Mini' crate with 4 slots (see Fig. 1) and the standard MPOD crate with 10 slots for HV modules. The most commonly used HV modules at GSI are 16-channel modules providing a positive or negative voltage up to 3000V at a maximum current of 3 mA.

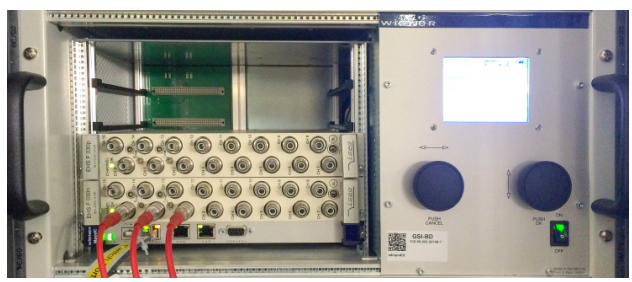


Figure 1: MPOD mini crate with four HV module slots. The lower two slots are populated with iseg modules.

CAEN features a different communication approach: The company provides proprietary libraries to interface with the hardware via Ethernet. The libraries are available for Windows and Linux and can be easily integrated into custom applications like FESA classes. The library offers an efficient monitoring feature for the current line of crates (SY4527 and SY5527). The so-called 'event mode' allows subscriptions to all available crate, board and channel properties. Whenever one of the subscribed values like current, voltage or status registers change, the hardware notifies the application of the change via the library. This way, the devices don't need to

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

UPDATE OF POWER SUPPLY CONTROL SYSTEM AT THE SAGA LIGHT SOURCE STORAGE RING

Y. Iwasaki#, T. Kaneyasu, Y. Takabayashi, S. Koda, SAGA Light Source, Tosu, Japan

Abstract

The control system for the SAGA Light Source storage ring power supplies is being upgraded to increase the ramp-up speed and allow the stored beam energy to be easily changed. By replacing the CPU module in the PLC used to control the power supplies, the ramp-up time was reduced from 4 to 2 minutes in a test bench prepared for the upgraded system. Until now, the allowable beam energy has been restricted to certain fixed values during ramp-up operation due to the original specifications of the PLC ladder program. To operate the storage ring at an arbitrary energy, the algorithm used in the PLC program was changed. The resolution at which the energy can be changed is less than 1 MeV. The upper layer of the control system, which uses National Instrument LabVIEW and ActiveX CA, was also reconstructed to produce a flexible GUI. Preliminary measurements of the energy dependence of the beam size and lifetime were carried out using the updated control system.

INTRODUCTION

The SAGA Light Source (SAGA-LS) is a middle-scale synchrotron radiation facility, at which electrons from a 255 MeV injector linear accelerator are injected into a 1.4 GeV storage ring [1]. The electrons are ramped up to 1.4 GeV in the storage ring during a period of 4 minutes. The storage current is 300 mA. The storage ring lattice is a double-bend type with 8-fold symmetry (see Figure 1). The power supplies for the storage ring magnets are listed in Table 1.

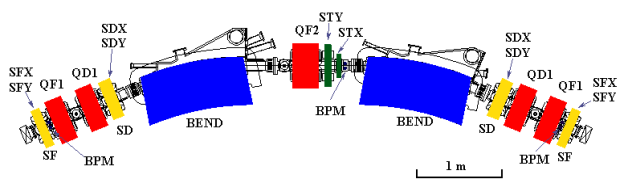


Figure 1: Unit cell of SAGA-LS storage ring.

The control system for the power supplies is constructed using personal computer (PC) and programmable logic controller (PLC) [2]. National Instruments LabVIEW is used for the graphical user interface (GUI). ActiveX CA [3], which provides EPICS channel access for Microsoft Windows, is adopted as the data communication protocol between server and client PCs.

#iwasaki@saga-ls.jp

Although the original control system worked stably, two problems existed with it. The first was the ramp-up speed. The ramp-up pattern file consists of 10,000 elements for each main power supply, and is stored in the internal register of a CUP module in the PLC. The ramp-up speed was limited by the performance of the CPU module. The second problem was a restricted accelerated beam energy. The allowable energy was fixed at certain values due to the original specifications of the PLC ladder program. We therefore replaced the CPU module: F3SP58-6S (YOKOGAWA Electric Corporation) with F3SP76-7S, and modified the algorithm of the ladder program slightly to allow arbitrary energy settings. The upper layer of the control system using LabVIEW and ActiveX CA was also reconstructed. Preliminary measurements of the energy dependence of the beam size and lifetime were performed using the updated control system.

In this paper, we describe the design of the control system for the power supplies at the SAGA-LS storage ring, the control algorithm and GUI, and the preliminary experimental results.

Table 1: Power Supplies for SAGA-LS Storage Ring Magnets

Power Supply	Control Interface	Number
Bending	16bit Digital	1
Quadrupole	16bit Digital	7
Sextupole	16bit Digital	2
BM-Corrector	Analogue , Digital	16
Q-Corrector	Analogue , Digital	40
Steering	Analogue , Digital	80

CONTROL LAYER FOR POWER SUPPLIES

A schematic view of the control layer for the power supplies is shown in Figure 2. Local control of the power supplies is performed by PLC. The PLC consists of one CPU module, an Ethernet module, I/O modules, and optical link bus modules. One CPU module controls all the power supplies of the storage ring. Sub-modules of the PLC are connected to the main unit with optical link bus modules. At the SAGA-LS storage ring, synchronization among power supplies is important, since the electrons are ramped up from 255 MeV to 1.4 GeV in a short time. The PLC outputs 10,000 set points for each

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

TESTING FRAMEWORK FOR THE LHC BEAM-BASED FEEDBACK SYSTEM

S. Jackson, D. Alves, L. Di Giulio, K. Fuchsberger, B. Kolad, J. Pedersen
CERN, Geneva, Switzerland

Abstract

During the first Large Hadron Collider (LHC) [1] shut-down period, software for the LHC Beam-based Feedback Controller (BFC) and accompanying Service Unit (BFSU) [2] was migrated to new 64-bit multi-core hardware and to a new version of CERN's FESA [3] real-time framework. This coincided with the transfer of responsibility to a new software team, charged with readying the systems for beam in 2015 as well as maintaining and improving the code-base in the future.

In order to facilitate the comprehension of the system's 90'000+ existing lines of code, a new testing framework was developed which would not only serve to define the system's functional specification, but also provide acceptance tests for future releases. This paper presents how the BFC and BFSU systems were decoupled from each other as well as from the LHC plant's measurement and actuator systems, thus allowing simulation-data driven instances to be deployed in a test environment. It also describes the resulting Java-based Domain-Specific Language (DSL) which allows the formation of repeatable acceptance tests.

INTRODUCTION

LHC operators rely on a feedback system in order to stabilise/correct and adjust the beams' closed-orbit, betatron tune, energy and radial loop [4] during the various stages of the operational cycle from filling to flat-top. The system is triggered at up to 25Hz by input signals from over 1'000 Beam Position Monitors (BPM), along with 6 tune measurements from Base-Band Q (BBQ) [5] measurement systems. The BFC sanitises this data, before calculating the necessary currents to send to various LHC steering dipoles and quadrupoles - The effects of which should be observed in the next iteration of the feedback loop (via the BFSU 1Hz instrumentation layer).

When a new team took charge of the existing codebase, it was apparent that many changes would be required to port the code to new 64 bit hardware, and to migrate to the latest FESA framework. This meant that testing the new software before the LHC start-up would be imperative. Testing using real signals from over 1'500 I/O devices round the LHC would not only be difficult during the machine's shutdown, but would be impossible when the LHC became operational. Consequently it was decided to create a testing framework which could emulate the BFC's input signals, the effect of which could then be observed via the BFSU at 1Hz. Test cases would then be written, for asserting that the BFC and BFSU were reacting correctly to the conditions defined by the test. It was also acknowledged that test developers may not be software experts, and may therefore be uncomfortable learning how to use a potentially complex API. In this context, a descriptive DSL would abstract the testing framework to be easily used by test creators.

Realising the framework, would not only require new testing code, but would entail many changes to the BFC and BFSU themselves, in order to allow instances of the software to be deployed outside their specialised operational hardware. Also, spoofing the BFC's input data would require relaxing of the BFC's data integrity checking mechanisms without compromising security in the operational system.

ADAPTING TO LIFE IN THE LAB

During 2014, the BFC and BFSU were ported to the new 64 bit architecture / FESA3 framework and the team had a release candidate ready for testing. Before any testing framework could be developed however, several changes to the software would be required to allow the software to be tested.

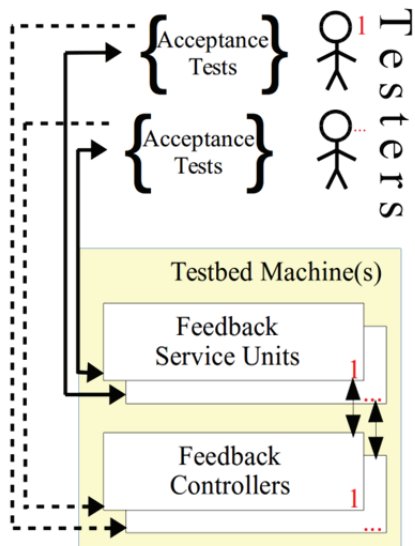


Figure 1: Software adapted so any number of BFC / BFSU pairs can be started, even on the same machine.

The new operational feedback system comprises of 2 multi-core Linux based (with Red Hat's MRG* kernel extension) HP Gen8 Blade machines. In order to assure real-time performance, the BFC is isolated from the LHC control system, with the BFSU acting as a client-facing data / setting proxy. A 2nd Ethernet link is used to transfer (~130Mb / sec) data between the BFC and BFSU

* Messaging Real-time Grid.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

ENHANCING THE DETECTOR CONTROL SYSTEM OF THE CMS EXPERIMENT WITH OBJECT ORIENTED MODELLING

R. Jiménez Estupiñán, A. Andronidis, T. Bawej, O. Chaze, C. Deldicque, M. Dobson, A. Dupont, D. Gigi, F. Glege, J. Hegeman, M. Janulis, L. Masetti, F. Meijers, E. Meschi, S. Morovic, C. Nunez-Barranco-Fernandez, L. Orsini, A. Petrucci, A. Racz, P. Roberts, H. Sakulin, C. Schwick, B. Stieger, S. Zaza (CERN, Geneva, Switzerland), U. Behrens (DESY, Hamburg, Germany), O. Holme (ETH Zurich, Switzerland), J. Andre, R. K. Mommsen, V. O'Dell (Fermilab, Batavia, Illinois, USA), Petr Zejdl (Fermilab, Batavia, Illinois; CERN, Geneva), G. Darlea, G. Gomez-Ceballos, C. Paus, K. Sumorok, J. Veverka (MIT, Cambridge, Massachusetts, USA), S. Erhan (UCLA, Los Angeles, California, USA), J. Branson, S. Cittolin, A. Holzner, M. Pieri (UCSD, La Jolla, California, USA)

Abstract

WinCC Open Architecture (WinCC OA) is used at CERN as the solution for many control system developments. This product models the process variables in structures known as *datapoints* and offers a custom procedural scripting language, called Control Language (CTRL). CTRL is also the language to program functionality of the native user interfaces (UI) and is used by the WinCC OA based CERN control system frameworks. CTRL does not support object oriented (OO) modelling by default. A lower level OO application programming interface (API) is provided, but requires significantly more expertise and development effort than CTRL. The Detector Control System group of the CMS experiment has developed CMSfwClass, a programming toolkit which adds OO behaviour to the *datapoints* and CTRL. CMSfwClass reduces the semantic gap between high level software design and the application domain. It increases maintainability, encapsulation, reusability and abstraction. This paper presents the details of the implementation as well as the benefits and use cases of CMSfwClass.

INTRODUCTION

The Detector Control System (DCS) applications of the CMS experiment are written with WinCC OA, mostly using the native CTRL language with the process variables being modelled in tree-like data structures, called *datapoints*. *Datapoints* are the default persistence layer and they are used by the runtime database to hold the process variables. CTRL language does not include a type definition syntax to declare and manipulate new structures. This can only be done by means of functions from the standard library or through the WinCC OA graphical editing interface. Low encapsulation and coupling between CTRL language and WinCC OA *datapoints* make data manipulation complex, compared to other languages.

The design of software, from the software engineering point of view, must be an independent process, unaware of the technology specifics [1]. During the formalization of a software design, abstract domain problems need to be transferred into algorithms and organized in data models. For that reason, the resources available in the programming platform, such as data structure definition mechanisms or programming language features, have a direct impact on the project implementation. When the semantic distance between the modelling language (e.g. UML) and the programming language is large, then the time and complexity of the translation process increases. Also other aspects of the software may be affected, like the readability of the code and maintainability of the entire software.

The CMS DCS team has attempted to close this gap by creating a development toolkit to add OO behaviour as well as code and data encapsulation down to the datapoint level. CMSfwClass toolkit revises the original WinCC OA and JCOP framework [2] device modelling concepts, enabling more abstract and powerful software architectural designs.

CMSFWCLASS TOOLKIT

The toolkit is composed of two different layers. The first layer is the back-end to add object orientation into WinCC OA. The second toolkit layer is a graphical user interface to provide easy and comprehensive access to the object oriented abstraction layer.



Figure 1: CMSfwClass toolkit modules.

Though the OO back-end has been designed to work in any WinCC OA project, the *GUI* has a certain level of integration with the CERN JCOP framework. Classes can be registered and accessed as JCOP device definitions. The *GUI* also uses CERN made libraries for syntax checking capabilities.

CMSfwClass OO Back-end Features

The OO back-end of CMSfwClass provides most of the commonly used features of modern OO programming and some custom object management features:

- Single inheritance and method overriding.
- Interface definition.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

LASER BEAM PROFILING AND FURTHER IMPROVEMENTS TO THE FHI FEL

H. Junkes, W. Schöllkopf, M. Wesemann, FHI, Berlin, Germany

Abstract

The free-electron laser facility at the Fritz-Haber-Institut (FHI FEL) startet regular operation in November 2013. Currently, the FEL provides continuously tunable, powerful pulsed laser radiation at any wavelength between 3.6 and 50 µm covering the mid-infrared (MIR) wavelength range. The first 1½ years of operation have opened up new possibilities for research in gas phase spectroscopy of molecules and clusters, non-linear optics, and surface chemistry at the FHI. The EPICS software framework was chosen to build the control system of this facility. Actions have been taken for laser beam profiling and characterization as well as improvements to the infrastructure like the archiver system. This paper presents the next upgrade step and lessons learned during the first time of operation.

DESIGN OF THE FHI FEL

The FHI FEL includes two oscillator FELs; a MIR branch for wavelengths up to about 50 µm and a far-infrared (FIR) branch for wavelengths from, possibly, about 40 to 400 µm. The MIR FEL has been commissioned and is fully operational whereas the FIR FEL has been projected as a future upgrade. A normal-conducting linear accelerator provides electrons of up to 50 MeV energy with a beam transport system feeding either of the two FEL branches or the electron diagnostics beamline as illustrated in Fig. 1 [1].

Electron Accelerator

The accelerator system was designed and build by Advanced Energy Systems, Inc.. It combines a thermionic electron gun, a sub-harmonic buncher cavity, and two S-band (2.99 GHz) standing-wave copper linacs. The first of the two linacs accelerates the electron bunches to a constant energy of 20 MeV. The second linac accelerates or decelerates the electrons to a final energy between 15 and 50 MeV. The accelerated electron bunches from bunch trains (macro-bunches) containing thousands of micro-bunches at a repetition rate of 1 GHz. The micro-bunch length can be compressed down to a minimum of 1 ps rms by a chicane between the linacs. The macro-bunches are repeated at 5 or 10 Hz [2].

Undulator and Cavity

In the MIR FEL a planar hybrid-magnet undulator is located within the IR cavity. The undulator is 2 m long containing 50 periods of 40 mm. It employs permanent magnets made out of NdFeB. At a minimum gap of nominally 16.5 mm, a maximum root-mean-square undulator parameter K_{rms} of more than 1.6 is reached. This, in combination with the minimum electron energy of 15 MeV, corresponds to a theoretical maximum wavelength of more than 50 µm for the MIR system.

The 5.4 m long MIR FEL cavity is formed by an end mirror and an out-coupling mirror. These are gold-plated copper mirrors of concave spherical shape. The waist of the cavity mode is located at the undulator center. A motorized in-vacuum mirror changer permits to select one out of 5 out-coupling mirrors with different outcoupling-hole diameters of 0.75, 1.0, 1.5, 2.5, and 3.5 mm for optimized performance at different wavelength regions. In addition, the cavity end mirror is mounted on a precision translation stage allowing for fine adjustment of the cavity length with 1 µm repeatability [2].

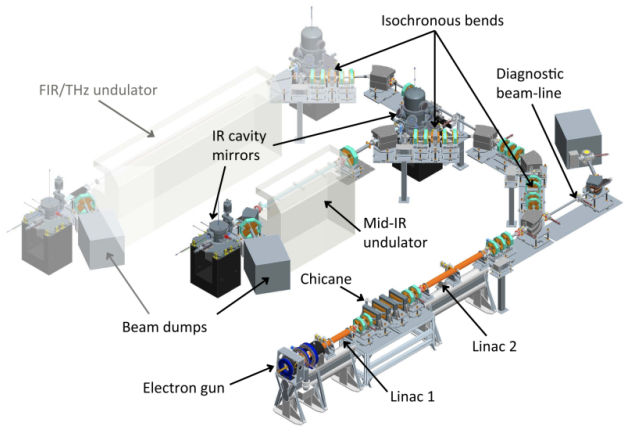


Figure 1: Overview of the FHI FEL installation.

LASER BEAM PROFILING

The IR pulses coupled-out from the FEL cavity pass through the evacuated IR beamline and propagate a distance of 18 m from the FEL vault to the diagnostic station located in the lab building. There the spectrum of the IR radiation is measured by an in vacuum Czerny-Turner grating spectrometer allowing for online monitoring of the FEL spectrum for each individual macro pulse. In addition, various commercial IR detectors are in use to determine the intensity of the FEL-pulses at different levels of sensitivity and temporal resolutions [2].

Fast IR Detector

A very fast IR detector (VIGO PEM-10.6) with sub 250 ps time resolution and an active area of 2 mm x 2 mm is used to detect individual micro-pulses. Figure 2 shows a measurement with this detector used in combination with a fast (4 GHz) oscilloscope without additional electrical amplification. At the standard micro-pulse repetition rate of 1 GHz individual micro-pulses are clearly resolved with a modulation depth of more than 50%. As can be seen in Fig. 2 this detector allows observing intensity changes on the single micro-pulse level.

Pyroelectric Linear Array

A pyroelectric linear array detector with 128 elements and a total width of 1/2" (DIAS Infrared 128LT) is

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

REAL-TIME ETHERCAT DRIVER FOR EPICS AND EMBEDDED LINUX AT PAUL SCHERRER INSTITUTE (PSI)

Dragutin Maier-Manojlovic, Paul Scherrer Institute (PSI), Villigen, Switzerland

Abstract

EtherCAT [1] bus and interface are widely used for external module and device control in accelerator environments at PSI, ranging from undulator communication, over basic I/O control to Machine Protection System for the new SwissFEL accelerator. A new combined EPICS/Linux driver has been developed at PSI, to allow for simple and mostly automatic setup of various EtherCAT configurations.

The new driver is capable of automatic scanning of the existing device and module layout, followed by self-configuration and finally autonomous operation of the EtherCAT bus real-time loop. If additional configuration is needed, the driver offers both user- and kernel-space APIs, as well as the command line interface for fast configuration or reading/writing the module entries.

The EtherCAT modules and their data objects (entries) are completely exposed by the driver, with each entry corresponding to a virtual file in the Linux *procfs* file system. This way, any user application can read or write the EtherCAT entries in a simple manner, even without using any of the supplied APIs. Finally, the driver offers EPICS [2] interface with automatic template generation from the scanned EtherCAT configuration. In this paper we describe the structure and techniques used to create the EtherCAT software support package at PSI.

INTRODUCTION

To support external device data acquisition and equipment control both for existing research facilities, such as Swiss Light Source (SLS), and for facilities being built at the time this text was written, like Swiss Free Electron Laser (SwissFEL) [3], an EtherCAT software interface was needed at PSI. Unfortunately, none of the existing commercial and non-commercial solutions we have reviewed and tested was able to cover and satisfy all of the requirements for the EtherCAT support.

General requirements were divided in two broad categories – the first was the full support for EPICS control system and the complete range of standard EPICS record types currently available, with the possibility for flexible addressing of EtherCAT modules and entries. The second requirement was to have a system that can provide the EtherCAT interface for common applications, both applications running locally, normally on the Ioxos IFC 1210 VME Board (equipped with the PowerPC P2020 CPU) and remotely, on a standard desktop PC or mobile device running any operating system capable of supporting network-based file systems, such as Linux, UNIX, Windows, MacOS, FreeBSD and others.

CONCEPTS

Providing support for such a wide range of applications in a single package presented a problem, since not every requirement or possible scenario for usage could have been satisfied with a single piece of software.

EPICS control system support requires its own type of dedicated device support driver. Unlike its kernel counterparts, EPICS driver has to run in Linux userspace, since EPICS system itself is a userspace application. Aside from EPICS, the system has to support other types of applications, both local and remote.

Local applications can be both userspace and kernelspace applications, which in turns mean at least two separate local APIs had to be created. Remote applications, on the other hand needed a generalized way to access EtherCAT data regardless of the operating system used.

EtherCAT Data Addressing

To describe an address of a given EtherCAT data entry, several variables have to be included: EtherCAT Master number (since there can be multiple masters running on the same host), Domain Nr. (domain is an arbitrary, user-defined collections of PDO (Process Data Object) entries sharing the same domain buffer memory and TCP exchange frame rate), Slave Nr. (Slave is another name for an EtherCAT Module), Synchronization Manager Nr. (SyncManager or SMs group objects by their exchange direction (input/output) or other criteria), Process Data Object Nr. (PDOs group entries by some arbitrary purpose defined by the Module producer) or Process Data Object Entries (PDO Entries or PDOE hold the actual data).

The user should be able to easily describe which data entry (or entries) should be addressed, in a consistent yet simple manner. To solve this problem, we have devised a new addressing schema for EtherCAT data:

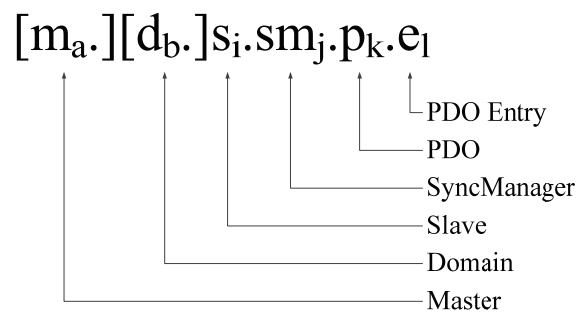


Figure 1: New schema for EtherCAT data addressing.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

PERSONNEL PROTECTION SYSTEM UPGRADE FOR THE LCLS ELECTRON BEAM LINAC*

M. Cyterski#, E. Chin,
SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

Abstract

As facilities age and evolve, constant effort is needed in upgrading control system infrastructure. This applies to all aspects of an accelerator facility. Portions of the Personnel Protection System of the Linac Coherent Light Source are still relying on a 50 year old relay based Safety System. This presents a substantial risk to the facility's ability to reliably perform its mission. An upgrade is underway to modernize these systems using Siemens S7-300 Safety PLCs and Pilz PNOZMulti programmable controllers. The upgrade will be rolled out over multiple years requiring the implementation to be fully compatible with adjacent legacy system while setting the foundation for the new generation system. The solution is a modularized safety system which can be deployed in a short time (1 month) while being flexible enough to adapt to the evolving needs over the next 20 years. Once fully deployed, the upgraded PPS System will provide not only greater availability to users, but also a higher level of Personnel Safety than previously provided.

BACKGROUND

The 3 mile Linac has been in operation at SLAC National Accelerator Lab for more than 50 years. In that time many upgrades have been made to support successful operation of the facility into the future. In 2014, a multiyear project was begun to upgrade the Personnel Protection System for the section of Linac used for operation of the Linac Coherent Light Source (LCLS).

The Personnel Protection System (PPS) at SLAC is composed of a system of interlocks whose function is to prevent personnel from being exposed to radiation within the accelerator enclosures.[1] In the Linac this is done by controlling the ability of the Klystrons to produce RF for the accelerator waveguides. The SLAC Linac is composed of 30-100meter long sections. Pairs of these sectors are configured as 15 stand-alone PPS zone.

The system currently used is composed of telephone relays and has operated safely for many years. In order to continue to support the high level of facility uptime demanded by LCLS, a modern replacement is needed. The LCLS Linac is composed of the final 10 sectors. The new system will provide a substantial increase in diagnostic information available to the operations staff as well as new functional modes. These modes will reduce facility downtime following personnel access to the accelerator housing.

*Work is supported by the U.S. Department of Energy, Office of Science under Contract DE-AC02-76SF00515.
#cyterski@slac.stanford.edu

Completing an upgrade of this size on an operational facility presents an additional level of difficulty. Annual maintenance down time is on the order of two months, with half of this time needed for recertification of all PPS installations. This only leaves four weeks to complete the removal, installation, and commissioning of a new system.

REQUIREMENTS

In addition to the standard requirements of a Personnel Safety System, this upgrade must meet two additional project specifics requirements. [2][3]

The first is the need for a personnel controlled access. Currently any access into the linear accelerator requires the operations staff to manually go into the tunnel and verify that no one is inside before closing the facility for beam operation. Depending on the number of sector pairs accessed this can take a substantial amount of time. Controlled Access will allow for trained personnel to enter the accelerator housing using a safety token and return the token once complete.[4] This eliminates the need for an operations search.

Secondly, due to the installation window time constraints, the upgrade must be deployed in phases. Each phase will be staged and modularly assembled while the facility is in operation and then deployed during a maintenance downtime.

DESIGN

The new system is composed of 3 separate PPS zones instead of the 5 sector pairs it's replacing. Each zone is deployed in a different annual maintenance down time. The three stages are listed in Table 1.

Table 1: Installation Stages

Zone	Controlled Access	Installation
Sectors 24 & 25	No	September 2014
Sectors 21-23	Sectors 21-25	September 2015
Sectors 26-30	Sectors 21-30	TBD 2016

Personnel Access to the accelerator regions is handled by three distinct access states at SLAC. No Access is the condition in which beam operation is allowed as long as the interlocks are met. Permitted Access allows free access to individuals to the PPS zone. This always requires an operations search of the region prior to machine recovery. Controlled Access allows for access without the need to search the PPS zone.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

UPGRADE OF THE CONTROL AND INTERLOCK SYSTEMS FOR THE MAGNET POWER SUPPLIES IN THE T2K PRIMARY BEAMLINE

Kazuo Nakayoshi, Ken Sakashita and Yoshiaki Fujii
High Energy Accelerator Research Organization (KEK), Tsukuba, Ibaraki, Japan

Abstract

T2K is a long-baseline neutrino oscillation experiment at J-PARC in Japan. A high intensity neutrino/antineutrino beam is generated and propagates 295km to Super-Kamiokande. The high intensity proton beam which reached 350 kW in May 2015, is extracted from the Main Ring synchrotron and guided through a primary proton beamline to a graphite target using normal-conducting (NC) magnets and super-conducting combined-function magnets. In October 2014, we replaced all of the power supplies (PSs) for the NC magnets with newly developed PSs. We also developed a new control system based on the Experimental Physics and Industrial Control System (EPICS) and PLC, putting emphasis on the safe operation of power supplies, and integrated it into the existing interlock system. Consequently the latency time for the interlock system was improved. We report the actual implementation and operation results of these developments.

INTRODUCTION

The T2K (Tokai-to-Kamioka) experiment [1] is a long-baseline neutrino oscillation experiment at J-PARC (Japan Proton Accelerator Research Complex) in Japan. A high intensity neutrinos/anti-neutrino beam is produced, and propagates 295 km from J-PARC to Super-Kamiokande.

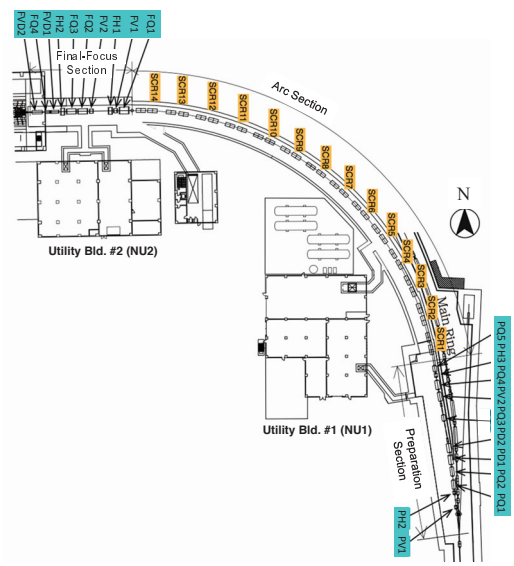


Figure 1: T2K primary beamline.

The T2K neutrino experimental facility is composed of primary/secondary beamlines and a near detector (ND280). The high intensity proton beam which reached 350 kW in May 2015, is extracted from the Main Ring synchrotron (MR) and guided through the neutrino primary beamline to a graphite target. Figure 1 shows the T2K primary beamline, consisting of the preparation (PREP), arc and the final focusing (FF) sections. Fourteen doublets of super-conducting combined function magnets [2] are located in the ARC section and bend the beam towards Kamioka. In the PREP section, the extracted proton beam is tuned with a series of 11 normal-conducting (NC) magnets. Ten NC magnets in the FF section guide and focus the beam onto the target. The power supplies (PSs) for the PREP section are located at the utility building-1 (NU1) and those for the FF section are located at the utility building-2 (NU2).

Some critical primary beamline issues are:

- A single failed beam shot, caused by a trip of the NC PSs, for example, results in serious damage of beamline equipments.
- The NC magnet power supplies were made mostly in 80's and required increasing effort for maintenance.

In order to solve these problems, we developed new power supplies for the NC magnets with a power supply company, and replaced all of the PSs in the summer of 2014.

NEW POWER SUPPLIES

The requirements for the new PSs for the NC magnets were (1) improvement of the safety interlock, (2) precise and stable operation, (3) improvement of maintenance, (4) upgrade of the control system and (5) downsizing. We report on each requirement and its actual implementation.

Table 1 shows the types of PS, corresponding to each magnet, its rated current, and its unit count. Figure 3 shows pictures of the new PSs.

Table 1: New power supplies for the normal conducting magnets in the T2K primary beamline.

Magnet type	DC OUT (A) / (V)	Converter type	Current stability(A)	# of units
Dipole	1500 / 100	chopper	0.1	4
Quadrupole	1000 / 100	chopper	0.1	9
Steering I	± 400 / ± 40	chopper	0.05	1
Steering II	± 200 / ± 20	switching	0.05	2
Steering III	± 100 / ± 10	switching	0.05	5

The output current stability is 0.1 A for the dipole and quadrupole magnet PSs and 0.05 A for the steering magnet PSs. The stability is superior to the previous PSs and the

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

INSTALLATION OF A HOT-SWAPPABLE SPARE INJECTOR LASER SYSTEM FOR THE SLAC LINAC COHERENT LIGHT SOURCE*

S. Alverson#, G. Brown, F.-J. Decker, S. Gilevich, S. Vetter, C. Zimmer
SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

Abstract

LCLS is a facility for generation of very short duration, highly intense x-ray pulses which requires an extremely reliable photocathode electron source. In order to maintain high uptime (>95%) for the experimenters, operations rely on a maintenance program for active laser components as well as on built-in redundancy in case of failure. To accomplish this, a duplicate laser system was installed, allowing for quick swap between the active system and the spare in the event of a malfunction or for planned maintenance. As an added bonus, this redundant system provides additional possibilities for science as both laser systems can also be run to the cathode simultaneously to create multiple particle bunches. Diagnostics were put in place to maintain both spatial and temporal overlap and allow for the fast switching between systems by operations personnel while still remaining within the safety envelope. This was done for both the primary UV drive laser as well as the secondary IR "heater" laser. This paper describes the installation challenges and design architecture for this backup laser system.

BACKGROUND

The Linac Coherent Light Source (LCLS) at SLAC National Accelerator Laboratory is a high intensity, extremely tuneable X-Ray Free Electron Laser Facility (XFEL). The light this facility uses for experiments is generated by a 1 km electron accelerator with a Photocathode source [1]. Light from a high power 253 nm wavelength (UV) laser (hereon referred to as the Drive Laser) is pulsed at up to 120 Hz on a copper cathode to emit electrons which are then accelerated by RF through a series of cavities.

In order to counteract unintended micro-bunching of the electrons in each pulse, they are also sent through a small chicane with an alternating series of permanent magnets (undulator) where they interact with a second 760 nm (IR) laser called the Heater Laser. To save the cost of constructing a completely separate system for the Heater Laser, it is actually formed from the unconverted light of the Drive Laser during the Tripling process used to create the UV light.

Loss of either of these laser systems due to an equipment failure would result in the unavailability of beam for the end users performing their experiment that week. Two exact duplicate laser systems (named Coherent 1 and Coherent 2) were developed so that there would always a spare available should something break or need to be serviced on the active line. Unfortunately, due to the high maintenance required to keep these lasers tuned up appropriately and the high sensitivity of the XFEL produced downstream to small changes in the characteristics of the injector laser systems, it could take many hours for a qualified laser specialist to switch to the spare, align it on the cathode, and tune it to match the old laser. This issue was compounded outside of the normal Day Shift hours when the specialist would have to drive in to site from home to perform the work.

REMOTE SWITCHING

In order to reduce the high downtime cost incurred by failures or scheduled maintenance tasks on the active laser system, it was desirable to implement a method by which Control Room staff could remotely switch between the lasers quickly. To do this, new mirrors and shutters were put in place in the common line of the Drive Laser and Heater Laser respectively before they entered the transport pipe down into the accelerator injector vault (Figure 1).

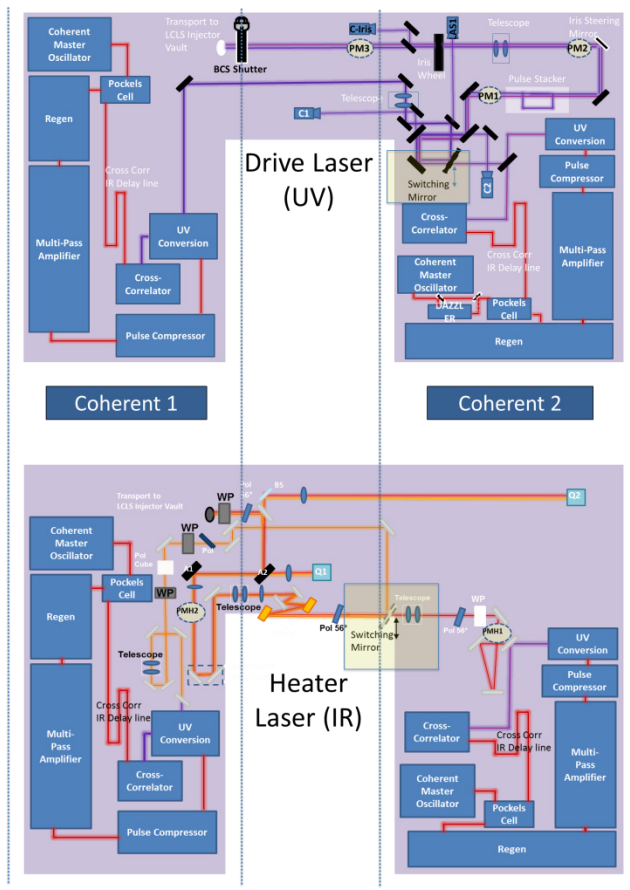


Figure 1: Optics table layout of active and spare injector laser systems.

* LCLS is supported by the U.S. Department of Energy, Office of Science under Contract DE-AC02-76SF00515.
#alverson@slac.stanford.edu

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

NEW DEVELOPMENTS ON EPICS DRIVERS, CLIENTS AND TOOLS AT SESAME

I. Saleh, Y.S. Dabain, A. Ismail, SESAME, Allan, Jordan

Abstract

SESAME is a 2.5 GeV synchrotron light source under construction in Allan, Jordan. The control system of SESAME is based on EPICS and CSS.

Various developments in EPICS drivers, clients, software tools and hardware have been done. This paper will present some of the main achievements: new linux-x86 EPICS drivers and soft IOCS developed for the Micro-Research Finland event timing system replacing the VME/VxWorks-based drivers; new EPICS drivers and clients developed for the Basler GigE cameras; an IOC deployment and management driver developed to monitor the numerous virtual machines running the soft IOCs, and to ease deployment of updates to these IOCs; an automated EPICS checking tool developed to aid in the review, validation and application of the in-house rules for all record databases; a new EPICS record type (mbbi2) developed to provide alarm features missing from the multibit binary records found in the base distribution of EPICS; a test of feasibility for replacing serial terminal servers with low-cost computers.

INTRODUCTION

SESAME consists of a 22 MeV Microtron, an 800 MeV Booster Synchrotron and a 2.5 GeV Storage Ring. Control System Implementation uses (EPICS) base R3.14.12. Servers are implemented as EPICS Input/ output Controllers (IOCs). Clients are implemented using a custom build of Control System Studio (CSS) based on V.3.16. Siemens S7 PLC controllers are used for the machine interlocks. An Allen Bradley PLC controller is used for the Personal Safety System (PSS). VME hardware is used for the timing system. Development and administration platforms use Scientific Linux 6.4. A Git version control is used to track development and documentation. All clients, servers, and controllers are connected to an isolated machine network. There are twelve virtual servers reserved to run the IOCs, archive system, alarm system and Git repositories.

The control systems have been implemented for the Microtron, Transfer Line 1 (TL1) and Booster. The Booster's control system is divided into seven subsystems: vacuum, power, RF, diagnostics, cooling, timing and Personal Safety System (PSS). Each control subsystem consists of one or more clients, servers, and controllers [1]. This paper will focus on the developments on EPICS drivers, clients and tools at SESAME.

TIMING SYSTEM DRIVER

The timing system for the Booster consists of one event generator VME-EVG230 and one event receiver VME-EVR230, both of which are connected to the EPICS network over TCP/IP. New Linux-x86 EPICS drivers were developed for the timing modules from Micro-Research Finland.

The VME-EVG230/VME-EVR230 are traditionally controlled over the VME bus. Both modules can also be programmed over the available Ethernet port however. Building an EPICS driver for controlling the timing modules over Ethernet instead of the VME bus has the following benefits:

- Drops the dead weight: The VME crate, the VME CPU card, the RTOS that runs on the CPU card (along with any required licenses), and the debug terminal that connects to the CPU card are no longer needed.
- Lowers the cost of implementation: A direct consequence of the point above.
- Confines the required development skills to Linux/x86 platforms. Knowledge in VME-bus, VxWorks, or any other RTOS/OS is not required.
- Maintains coherency in the control infrastructure. This point may be specific to SESAME only since all of the IOC's at SESAME run on Linux/x86 platforms.

The device layer uses the traditional asynchronous processing model of EPICS support modules. The driver layer uses the remote programming protocol of the timing modules to control them. This protocol uses UDP. The drivers add feedback, retransmissions, and timeouts to create reliable communication over UDP.

Both drivers are LGPL'd and can be found as public domain on Github [2]. The client of the timing system is implemented in Control System Studio (CSS) and is shown in (Fig. 1).

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

CONTROL SYSTEM UPGRADE FOR SUPERKEKB INJECTOR LINAC

M. Satoh#, F. Miyahara, Y. Seimiya, K. Mikawa, T. Suwada, K. Furukawa, KEK/ SOKENDAI, Tsukuba, Japan
T. Kudou, T. Ichikawa, Y. Mizukawa, K. Hisazumi, S. Kusano, Mitsubishi Electric System & Service Co., Ltd, Tsukuba, Japan
H. Saotome, M. Takagi, Kanto Information Service (KIS), Accelerator Group, Tsuchiura, Japan

Abstract

The KEKB project has successfully ended 13 years operation in the June of 2010. The construction of SuperKEKB main ring has almost completed for aiming at the peak luminosity of 8 x 10^{35} $cm^{-2}s^{-1}$. The injector linac upgrade is also going on for increasing the intensity of bunched charge with keeping the small emittance. The key upgrade issues are the stable operation of a new positron capture system, a low emittance photo-cathode rf electron source, and the establishment of low emittance preservation. The injector linac beam commissioning started in the October of 2013. The control system performance determines the beam operation efficiency of injector linac, and it eventually has a strong impact on the experimental results of physics. In this decade, the linac control system has been gradually transferred from the in-house system to the Experimental Physics and Industrial Control System based one for the high availability of beam operation. In this paper, we present the detail of SuperKEKB injector linac control system.

INTRODUCTION

The linac beam control system is based on a standard client and server model with three layers. It comprises a client, a server, and a local controller layers. At the beginning stage of KEKB operation, the original linac control system has been developed by using the in-house software libraries based on the remote procedure call (RPC). The client user interfaces have been implemented by a command line interface based on shell script and a graphical user interface based on Tcl/Tk scripting language. Around a decade ago, the middle phase of KEKB operation, the linac control system has been upgraded to a new framework based-on the Experimental Physics and Industrial Control System (EPICS) [1] to improve the affinity between the linac and KEKB main ring control systems. These improvements make it possible to conduct the correlation analysis of linac and main ring parameters in much easy way. Such analysis can strongly help to find a source of injection rate deterioration and any other troubles.

After introducing the EPICS framework, the new client side user interfaces have been implemented by Python scripting language since its rich library modules and functionalities can accelerate the software development process. For the simultaneous top-up injection of KEKB and PF rings, an event based timing system has been installed to enable the pulse-to-pulse beam modulation [2].

OUTLINE OF CONTROL SYSTEM

Computer Environment

In the beginning phase of KEKB project, the six Compaq Alpha servers based on the Tru64 UNIX operating system were utilized as the server computers. All server programs were basically run on them. Two of them were connected to the RAID disk drive via SCSI bus interface, and they can work as the active/standby redundant NFS servers for aiming at the high availability operation. Since the Tru64 UNIX was obsoleted, all server machines have been gradually replaced by the Linux-based servers. Eventually, we made decision to use the Linux base system without cluster functionality after the evaluation of several different types of high availability cluster system based on Linux. For keeping the high availability of server machines without cluster scheme, the blade server system was installed as the high reliability server. Currently, fifteen Linux based blade servers are utilized as the server computers with CentOS 5.11 x86_64. Both of server and client side control software are running on the server machines. CentOS 6.7 and 7.1 running on two blade servers are currently tested by using VMware vCenter of the virtual computing technology.

The Tektronix X terminals and touch panel displays based on PC9801/DOS machines have been originally utilized as the operator terminals. These terminals were replaced by the Linux based PCs and Windows based PCs with the X server application software of ASTEC-X and

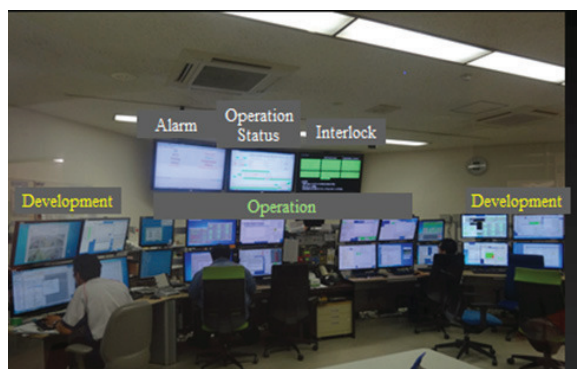


Figure 1: Photograph of the KEK electron/positron injector linac main control room.

e-mail address: masanori.satoh@kek.jp.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

CONTROL SYSTEM DEVELOPMENTS AT THE ELECTRON STORAGE RING DELTA

D. Schirmer*, A. Althaus, F. Bahnsen
Center for Synchrotron Radiation (DELTA), TU Dortmund, Germany

Abstract

Increasing demands, mandatory replacement of obsolete controls equipment as well as the introduction of new software and hardware technologies with short innovation cycles are some of the reasons why control systems need to be revised continuously. Thus, also at the EPICS-based DELTA control system [1], several projects have been tackled in recent years:

(1) Embedding the new CHG[1]-based short-pulse facility for VUV and THz radiation [2-4] required, for example, the integration of IP-cameras, Raspberry-Pi PCs and EtherCat/TwinCat wired I/O-devices [5, 6].

(2) The request for a staff-free control room led to the programming of new web applications using Python [7] and the Django framework [8]. This development resulted in a web-based interlock system that can be run, amongst others, on Android-based mobile devices.

(3) The virtualization infrastructure for server consolidation has been extended and migrated from XEN [9] to the kernel based KVM [10] approach.

(4) I/O-units which were connected via conventional fieldbus systems (CAN, GPIB, RS-232/485), are now gradually replaced by TCP/IP-controlled devices.

This paper describes details of these upgrades and further new developments.

INTRODUCTION

DELTA, a 1.5-GeV electron storage ring, is operated since 1999 by the TU Dortmund University as a synchrotron light source for campus-based, regional and international users. Since 2011, the facility has been extended by a short-pulse source for VUV and THz radiation making use of the CHG principle. An upgrade to EEHG[2] is in preparation [11]. Not only for these reasons the EPICS-based DELTA control system has been revised and complemented in many fields.

NEW SOFTWARE

Network Administration Tool

Since the number of control system network devices and, thus, the complexity of the network topology increased constantly, it was mandatory to develop a DELTA-specific management tool. With the help of this tool, all individual devices and their connection properties are registered centrally. It administers network interfaces, IP/MAC numbers and assignments to domains. Furthermore, it generates consistent DNS/DHCP configuration files as well as location and network plans. The program is implemented as a web application and is based on the high-level Django/Python framework [7, 8]. All data and configurations are stored in the DELTA MySQL [12] database.

Web Applications

The local radiation safety authority has given its consent to a non-permanent attendance of the control room during standard machine operation. For security reasons, however, it was necessary to develop an additional alert system, which notifies the operator to machine malfunctions via a mobile phone.

For real-time monitoring of the accelerator interlock warnings, a web-socket connection is performed, which keeps a full-duplex communication link between a client (e.g., an Android-based mobile phone) and a server continuously open. The connection uses WAMP[1] [13], a web-socket sub-protocol (with RPC[2] and publish & subscribe mechanism).

On the server side (see Fig. 1 below), *Crossbar.io* [14] is the central software component. It interacts as a message WAMP-router for different so-called service nodes. One node provides the service to read EPICS records another calculates und triggers interlock warnings (all implemented with Python). In addition, it launches web services that provide panels accessible from various URLs. There are web pages for machine status (e.g., beam current, lifetime, insertion device status), interlock notifications as well as the electronic logbook (see Fig. 2).

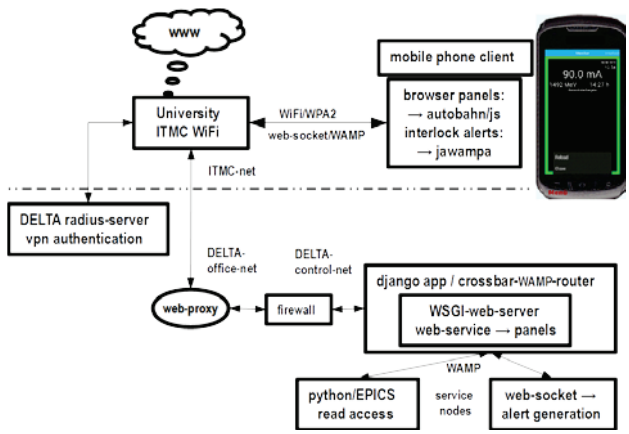


Figure 1: Software setup of the WiFi alert system.

On the client side (see Fig. 1 top), e.g., an Android-based smartphone, an *Autobahn/JS*-application [15] (an implementation of WAMP in javascript) displays the

*detlev.schirmer@tu-dortmund.de
[1]Coherent Harmonic Generation
[2]Echo -Enable Harmonic Generation

[1]Web Application Messaging Protocol
[2]Remote Procedure Call

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

UPGRADES TO CONTROL ROOM KNOBS AT SLAC NATIONAL ACCELERATOR LABORATORY *

S. Hoobler#, S. Alverson, M. Cyterski, R. Sass, SLAC National Accelerator Laboratory, Menlo Park, CA, 94025, U.S.A

Abstract

For years, accelerator operators at the SLAC National Accelerator Laboratory (SLAC) have favored hardware knobs in the control room for accelerator tuning. Hardware knobs provide a tactile, intuitive, and efficient means of adjusting devices. The evolution of separate control systems for different accelerator facilities at SLAC has resulted in multiple flavors of knob hardware and software. To improve efficiency, space usage, and ease of use, the knob systems have been upgraded and integrated.

BACKGROUND

SLC Knobs

In the 1990s a knob system was developed for the SLAC Linear Collider (SLC), using hardware knob chassis designed by Miranova Systems. SLAC developed software and a DOS knob controller to interface the knob chassis to the SLC VMS control system [1].

EPICS Knobs

The Linac Coherent Light Source control system is based on EPICS (Experimental Physics and Industrial Control System). EPICS software was written to interface the existing knob hardware to the LCLS control system. At that time, it was presumed that the VMS control system would soon be decommissioned and no effort was made to integrate the two systems.

Knob Use at SLAC

To date, the VMS control system and its knobs are still used by the Facility for Advanced Accelerator Experimental Tests (FACET) and the Next Linear Collider Test Accelerator (NLCTA). EPICS knobs are used at LCLS, the SPEAR storage ring, and NLCTA.

Until recently, the VMS and EPICS knob systems were still not integrated, making an inefficient use of space and knob hardware. To improve this, the knob software and hardware have been updated to allow individual boxes to be used by multiple control systems simultaneously.

DESIGN CONSIDERATIONS

Backward Compatibility

To stay compatible with the legacy knob hardware and software, the new system supports the existing communication protocols.

*Work supported by U.S. Department of Energy under Contract Number DE-AC02-76SF00515

System Performance

The system should maintain the performance standard of the previous systems.

- Good knob resolution.
- Quick response.
- Allow same device to be assigned to multiple boxes simultaneously without conflict.
- Ability to restore device to original value.

Enhancements

To support operation of multiple facilities from the same location, the system was enhanced to allow a single box to be used with multiple control systems simultaneously.

Maintainability

The system will initially use the legacy protocols to remain backward compatible. As older systems become decommissioned, it may be desirable to migrate to other protocols. The software should be modular to facilitate this.

Flexibility

Unlike the legacy SLC control system, the LCLS EPICS control system is distributed. There is no pre-defined list of devices that may be attached to knobs. Any EPICS process variable may be attached. The software should be flexible to handle generic cases and should also provide support for device-specific settings and facility-specific conventions.

SYSTEM IMPLEMENTATION

A knob system consists of the elements shown in Figure 1. The Knob Controller handles all communication with the hardware knobs. The hardware knob boxes and the Knob Controller together act as a server to the Knob Users. The Knob Users are the accelerator control systems which wish to attach devices to knobs. The Controller is the interface between the Users and the hardware knobs and will accept requests from any User in the system.

In this implementation, the knob hardware uses a serial protocol while all software entities use Ethernet protocols. A terminal server performs the conversion between serial and Ethernet packets. An individual knob box has four knobs, each of which can be independently controlled by a different control system. A Knob Controller and terminal server can support up to 16 knob boxes.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

DESIGN AND COMMISSIONING RESULTS OF MICROTCA STRIPLINE BPM SYSTEM*

S. Hoobler#, R. Larsen, H. Loos, J. Olsen, S. Smith, T. Straumann, A. Young, C. Xu, SLAC National Accelerator Laboratory, Menlo Park, CA, 94025, U.S.A.
H. Kang, C. Kim, S. Lee, G. Mun, Pohang Accelerator Laboratory, Pohang, Kyungbuk, Korea

Abstract

The Linac Coherent Light Source (LCLS) is a free electron laser (FEL) facility operating at the SLAC National Accelerator Laboratory (SLAC). A stripline beam position monitor (BPM) system was developed at SLAC [1] to meet the performance requirements necessary to provide high-quality stable beams for LCLS. This design has been modified to achieve improved position resolution in a more compact form factor. Prototype installations of this system have been operating in the LCLS LINAC and tested at the Pohang Accelerator Laboratory (PAL). Production systems are deployed at the new PAL XFEL facility and at the SPEAR storage ring at the Stanford Synchrotron Radiation Lightsource at SLAC. This paper presents the design and commissioning results of this system.

BACKGROUND

At LCLS, stripline BPMs measure the transverse position of the electron bunch in the injector, Linac, and transport lines. Each BPM is instrumented with a SLAC-designed chassis that houses analog conditioning electronics, a digitizer, and a CPU. After digitization, the processed data is transmitted over a dedicated network to a VME processor that calculates pulse-by-pulse beam position. An online calibration scheme injects a tone into the system between beam pulses in order to compensate for gain variation and drift. This process occurs at the beam rate, which is 120 Hz.

For recent projects, it became desirable to use electronics in a smaller form factor while maintaining the performance standard of the LCLS design.

SYSTEM DESIGN

A new system was developed, using the MicroTCA (Micro Telecommunication Computing Architecture) crate architecture. In this system, electronics for several BPMs are housed within a single crate, with digitizers located in the front of the crate and the analog conditioning electronics in the rear. Digitized data is transferred along the PCIe backplane to the CPU for processing.

The first version of this system [2] was tested in the LCLS Linac and found to meet the performance requirements. Two of these BPMs have been deployed in a SPEAR transport line.

Later, this design was further modified to achieve improved position resolution for the PAL XFEL facility. Revised electronics, including a different operating frequency, larger bandwidth, and faster digitization rate, provide position resolution improved by a factor of two. Prototypes of this system were tested in the SLAC Linac and at the PAL Injector Test Facility (ITF). Hardware design details presented in this paper are for this improved design.

Analog Front End

The analog front end (AFE) is used to reduce the signal bandwidth and set the signal level before digitization. To achieve the desired system resolution at low bunch charge, it is important to minimize losses in the system. This is most important in the signal cables and the AFE components up to the first amplifier.

The SLAC-designed AFE contains two stages of amplifiers and attenuators. The attenuators can be adjusted from the control system to maintain reasonable signal level when the bunch charge changes. The bandpass filters select a processing frequency of 300 MHz with a bandwidth of 30 MHz. Figure 1 shows a block diagram of a single input channel of the AFE.

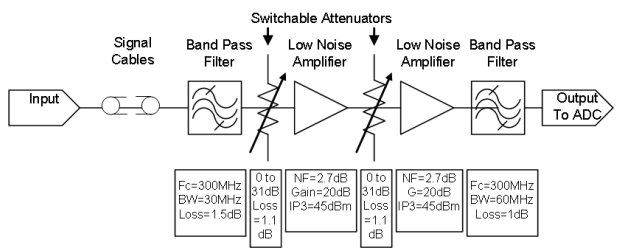


Figure 1: Analog Front End block diagram.

As in the LCLS system, the AFE performs a self-calibration process between beam pulses. A 300 MHz tone is injected into a single stripline and the coupled response is measured from the two adjacent striplines. This is performed for both the horizontal and vertical transverse planes. The calculated calibration ratios are applied to the beam position calculation to compensate for thermal drift and gain variation among channels. Figure 2 shows a block diagram of the calibration scheme.

*Work supported by U.S. Department of Energy under Contract Numbers DE-AC02-76SF00515 and WFOA13-197

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

TIP: AN UMBRELLA APPLICATION FOR ALL SCADA-BASED APPLICATIONS FOR THE CERN TECHNICAL INFRASTRUCTURE

Ph. Gayet, P. Golonka, M. Gonzalez-Berges, L. Goralczyk, J. Pache, P. Sollander, F. Varela[#], CERN, Geneva, Switzerland

Abstract

The SCADA package WinCC Open Architecture (WinCC OA [1]) and the control frameworks (JCOP [2], UNICOS [3]) were successfully used to implement many critical control systems at CERN. In the recent years, the Industrial Controls and Electronics (ICE) group of the Engineering Department (EN) at CERN, supported other groups to re-implement the supervision of technical infrastructure, like Electrical distribution (EL), and Cooling and Ventilation (CV), using these tools. However, the fact that these applications are highly independent from the operation point of view, as well as their increasing number, renders operation uncomfortable since shifters are forced to continuously switch between them. In order improve the integration, EN-ICE is developing the Technical Infrastructure Portal (TIP) to provide centralized access to all WinCC OA applications, and extending their functionality including links to external databases and to a powerful localization system based on GIS. In addition the tool offers an environment for operators to develop views that aggregate data from different sources, such as cooling and electricity. This paper also describes the challenges faced during the implementation of TIP due to the large degree of heterogeneity across applications, although they are made out of common building blocks.

INTRODUCTION

For the construction phase of the LHC, a big emphasis was put in the use of industrial tools for the control systems. This was complemented with the development of frameworks that provided the necessary adaptations to the CERN environment (JCOP, UNICOS). This approach showed to be very successful and many systems were built following it in both the experiments (e.g. Detector Control Systems, Detector Safety Systems, Gas Control, etc.) and the accelerators (e.g. Cryogenics, Machine Protection, Vacuum, etc.). The result was a big gain in manpower and a high homogenization across systems.

Building on this success, an effort was started in 2009 to use the same tools for the CERN infrastructure systems. These systems cover areas like Cooling and Ventilation, Electrical Distribution or Radiation Monitoring. There is now a large number of them that have benefited from this standardization effort, and there is an ongoing process to converge to the same technologies for most of the applications, if not all. The homogenization that results from this convergence process, opens the door to further integration at the SCADA layer and better operations efficiency. This paper shows the first initiative to build a hypervisor system on top of the existing applications across several domains for this purpose.

OPERATION OF THE CERN TECHNICAL INFRASTRUCTURE

Operators on shift in the CERN Control Center (CCC) are in charge of monitoring the Technical Infrastructure (TI) 24 hours a day, 365 days per year. Their duties are to detect problems, make first diagnostics and decide on actions to take. The TI encompasses the entire CERN electrical distribution network from 400 kV down to 48 V, the CV installations for all accelerators, detectors and other important facilities such as the computer centre and other systems (controlled access, safety monitoring systems, etc.). Anomalies are detected through a centralised alarm display, LASER [4], to which all systems across CERN are connected. Diagnostics are made using the application specific supervision systems.

There are more than 100 standalone CV systems around the CERN sites and their corresponding supervision systems have been developed over time with different technologies and by different people. This makes difficult for the operators to treat each non conformity as there is no standard way to establish a correlation between an alarm reported in the LASER console and the corresponding synoptic view of the process experiencing the problem. This requires that the operators have a very detailed knowledge of the various systems and makes the learning curve for newcomers very steep. Moreover, each of these applications requires individual login making daily life very tedious for the operators.

TI PORTAL

Main Requirements

TIP shall centralize the access to all WinCC OA applications deployed for the CERN TI simplifying the work of the operators in the control room. TIP must provide access to all synoptic views and trend plots available in every remote application, and must use the access control schemas locally defined in each of these applications:

- a single log-in shall be sufficient to grant the operator access to every single remote application,
- the operation can be performed by both the equipment group expert and the central operators with privilege associated to their level of expertise,
- the data shall be presented to the operators and experts in exactly the same way,

[#]fernando.varela.rodriguez@cern.ch

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

KECK TELESCOPE CONTROL SYSTEM UPGRADE

K. Tsubota, J. A. Mader, W. M. Keck Observatory, Kamuela, Hawaii

Abstract

The Keck telescopes, located at one of the world's premier sites for astronomy, were the first of a new generation of very large ground-based optical/infrared telescopes with the first Keck telescope beginning science operations in May of 1993, and the second in October of 1996. The components of the telescopes and control systems are more than 15 years old. The upgrade to the control systems of the telescopes consists of mechanical, electrical, software and network components with the overall goals of improving performance, increasing reliability, addressing serious obsolescence issues and providing a knowledge refresh. This paper is a continuation of one published at the 2013 conference [1] and will describe the current status of the control systems upgrade. It will detail the implementation and testing for the Keck II telescope, including successes and challenges met to date. Transitioning to night time operations will be discussed, as will implementation on the Keck I telescope. Throughout this paper the telescope control system will be referred to as TCS.

TCS OVERVIEW AND CONTEXT

Each of the Keck telescopes is over 8 stories tall, weighs in excess of 300 tons, uses an alt-az mount and has the equivalent of a 10meter mirror which is comprised of 36 individual segments. A general diagram is shown in Fig. 1. The telescope supports a wide range of instruments at various focal locations in addition to providing both a natural guide star and laser guide star adaptive optics.

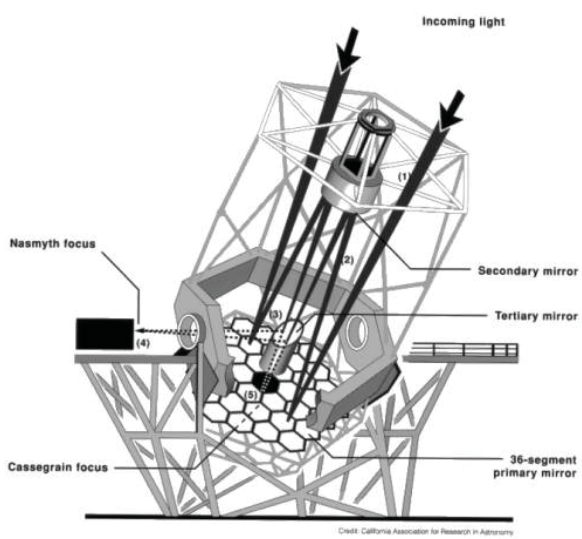


Figure 1: Keck Telescope.

TCS consists of a supervisor and set of subsystems that work together to provide Status and Control of:

- Telescope mounts
- Dome and shutter positions
- Facility rotators
- The secondary mirror

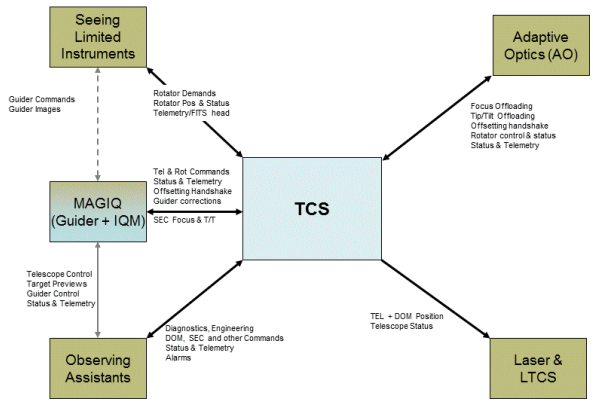


Figure 2: TCS System Context.

Figure 2 shows a high level context diagram for the Telescope Control System (TCS). It and the other Observatory principal systems interact through well-defined interfaces to accomplish the desired behaviour. Systems are tied together by the use of an Ethernet Bus. The main purpose of the TCS software is to accept the target position of a celestial object (which can be given in a variety of coordinate systems) and then calculate the mount, rotator and dome/shutter positions, so that the target is imaged perfectly at a given point in the focal plane. Furthermore, the TCS is characterized by the need to integrate a number of heterogeneous subsystems, which exhibit complex interactions. These interactions, although not hard real-time bounded, need a high level of synchronization. All this has to be done in a manner that protects the staff and equipment.

TCS operates the mechanical components of the telescope through a number of computer subsystems which control the telescope mount assembly, the dome enclosure, secondary and tertiary mirrors and facility rotators. The subsystems are responsible for the actual servo or hardware control, while the TCS coordinates their activities.

- The Telescope Axes (AXE) subsystem controls the telescope mount assembly, including the elevation and azimuth drives and other components associated with the mount operation.
- The Dome subsystem (DOM) operates the enclosure, which includes the dome carousel, dome lights and shutters, and other auxiliary equipment.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

EPICS IOC BASED ON COMPUTER-ON-MODULE FOR THE LNL LABORATORY

J. Vásquez, D. Pedretti, R. Ponchia, INFN/LNL, Legnaro, Italy
M. Bellato, R. Isocrate, INFN, Padova, Italy
M. Bertocco, UNIPD, Padova, Italy

Abstract

At LNL it is being carried out an upgrade campaign of the control systems of the accelerator complex. The two main goals are standardization of hardware and software and system interoperability. EPICS has been chosen as the standard framework for developing new control systems; this will address software standardization and system interoperability. In order to achieve hardware standardization, a new EPICS IOC is under development, which will become a basic construction block for all future control systems. The COM (Computer-on-Modules) from factor has been chosen as the hardware platform for the IOC, along with the peripheral devices needed for developing all the foreseen control system at LNL. Prototypes of this IOC has been developed using ADLINK's Type 6 COM Express modules on generic carrier boards with DIO, ADC and DAC expansion boards. These prototypes have been tested under typical applications at LNL in order to validate the hardware platform choice. Experimental results show that the performance of the IOC in terms of effective resolution (ENOB and bias error), sample rates and CPU usage is suitable for satisfying the requirements of the control systems.

IOC DESCRIPTION

The IOC is intended as a standard system with the necessary interfaces in order to be able to control the instrumentation present on all the foreseen system at LNL. EPICS tools will be used for software development in a standard way. For each application, a custom-made control algorithm will be developed under the framework.

COM form factor has been chosen as the hardware platform for the IOC. In particular, the type 6 COM Express standard was selected [1].

The COM will be installed on a custom carrier board, designed at LNL. This carrier board will contain peripheral devices such as digital IO, ADCs, DACs, stepper motor controllers, and communication interfaces, among others, connected to the COM. Its design is under carefully considerations in order to cover all the common needs for the current and the foreseen future control systems at LNL.

Some interfaces are already available directly from the COM. Others will be controlled using USB ports. Many others will be controlled by a FPGA that will communicate with the COM through a PIC Express lane. Figure 1 shows a simplify block diagram of the IOC.

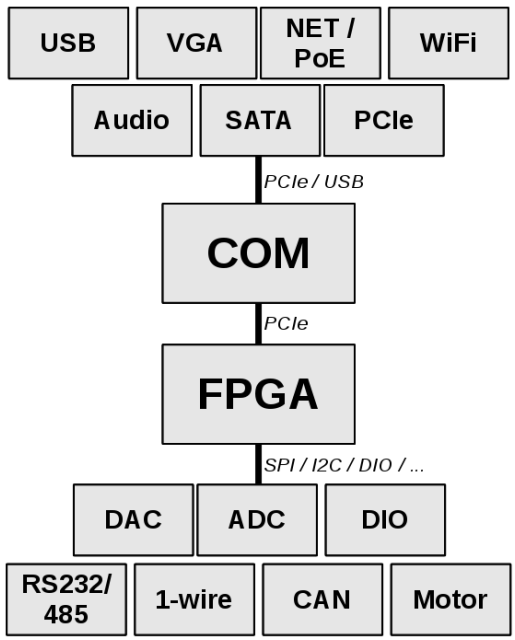


Figure 1: IOC simplify block.

IOC PROTOTYPES IMPLEMENTATION AT LNL

In order to test the validity of the COM as hardware platform for IOC developments, prototypes were developed using commercial devices. These prototypes, at the same time, allowed developing the software part of control systems.

The developed software is mostly independent from the hardware implementation, due the fact that the selected COM uses Intel CPUs allowing the use of a standard Linux environment for the implementation. In this way, once the custom IOC will be ready, the software would be easy transport to the new hardware platform.

On the other hand, prototypes were used to develop and implement diverse control systems at LNL. This allowed testing them under real operative conditions, as well as to develop critical control system without waiting for the final custom IOC to be available.

Hardware Architecture

The prototypes have been developed using Adlink [2] Type 6 COM Express module on a generic carrier board. Commercial Digital IO, ADCs and DACs PCIe boards were installed on the carrier board in order to provide IO capabilities to the IOC.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

MEBT AND D-Plate CONTROL SYSTEM STATUS OF THE LINEAR IFMIF PROTOTYPE ACCELERATOR*

J. Calvo†, D. Jiménez-Rey, E. Molina Marinas, J. Mollá, I. Podadera. Ciemat, Madrid. Spain

Abstract

Linear IFMIF[1] [1] Prototype Accelerator (LIPAc) [2], Rokkasho, Japan, comprises a succession of devices and systems that accelerate a deuteron beam up to 9 MeV with a current of 125 mA, generating a power of 1.125 MW, and transport it up to a beam dump. The beam power becomes critical from the point of view of losses; even tiny losses must be avoided. This fact, and the complexity of the accelerator operation, requires a coherent strategy when designing, commissioning and optimizing the accelerator control system, specifically focused in the control systems of the Medium Energy Beam Transport (MEBT) and the Diagnostic Plate (DP, a movable set of diagnostics). Both systems are essential to validate the performance of the accelerator and particularly the ion source, Radio Frequency (RF) and Radio Frequency Quadrupole (RFQ) systems. This contribution will describe the recent advances in the control architectures and the EPICS based developments achieved in MEBT for the motion control of bunchers and scrapers, control of the power supplies in quadrupoles and steerers, and refrigeration and vacuum. Besides, control of fluorescence profile monitors (FPMs) in the D-Plate is displayed.

LIPAC CONTROL SYSTEM

LIPAc control system consists of the remote control, monitoring and data acquisition of all devices, systems, subsystems and operations carried out in the accelerator vault with intense radioactive environments. It uses EPICS [3] as the main set of control software tools, providing the control with distribute and real-time features. It is composed of five main elements, namely, accelerator Central Control System (CSS), Local Area Network (LAN), Personal Protection System (PPS) [4], Machine Protection System (MPS) and Timing System (TS). Within each LIPAc subsystem there is a Local Control System (LCS).

LIPAc control system development is an unique and complex work due to the special characteristics of the prototype accelerator and its beam dynamics, that are, the unprecedented high beam intensity inducing the simultaneous combination of high beam power and high space charge. Therefore, MEBT and D-plate control system development is related with general accelerator characteristics in the following way:

- LIPAc is aiming to produce a very powerful CW deuteron beam (1.1 MW @ 9 MeV). The unavoidable beam losses [5] along the accelerator lead to an important rate of neutrons and radiations that must be properly controlled. On the other hand, the power of the beam itself could, in case of accidental or abnormal operation conditions, damage or destroy highly valuable parts of the machine.
- Beam losses must be observed and studied in order to fulfil beam stability in terms of transverse position and shape. To obtain this fundamental information several diagnostics devices have been installed along the accelerator line, these devices require an specific control and a fast data acquisition.

LIPAc Local Control Systems (LCSs)

The main function of the LCSs is, obviously, to control its own subsystem, displaying their own process variables (PVs) for the rest of LCSs and enabling access and modification of these PVs by the CSS. Safety, device protection and optimization activities are also included among its functions. LCSs are composed of different elements but following three layers architecture based on EPICS, Fig. 1. The lower level consists of equipment (power supplies, diagnostics, sensors, motors) connected to the intermediate level via specialized links. The middle layer consists of EPICS front-end systems (so-called input-output controllers (IOCs)). The upper layer, part of the CCS, consists of workstations, which either play dedicated roles, offering a man-machine interface function. These two levels communicate through TCP/IP on Ethernet.

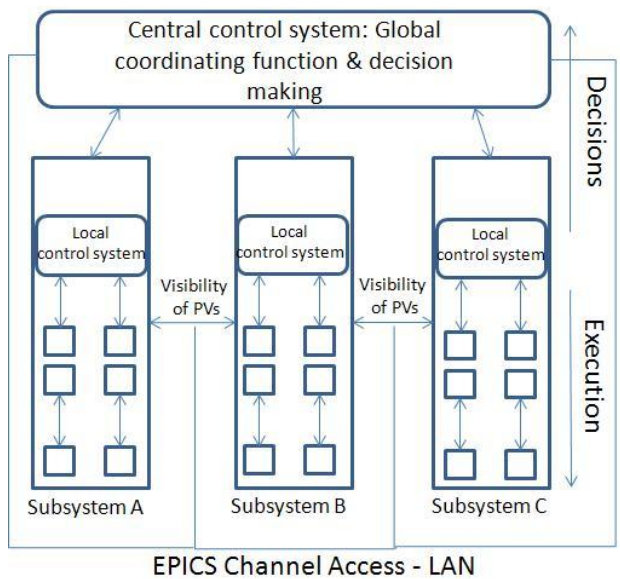


Figure 1: LIPAc general distributed process control.

* This work has been partially funded by the Spanish Ministry of Economy and Competitiveness, under projects OPTIMHAC FIS2013-40860-R and IFMIF-EVEDA II. Ref: AIC-A-2011-0654.

† julio.calvo@ciemat.es

[1] IFMIF, the International Fusion Materials Irradiation Facility, is an accelerator-based neutron source that will use Li (d, xn) reactions to generate a flux of neutrons with a broad peak at 14 MeV equivalent to the conditions of the Deuterium-Tritium reactions in a fusion power plant. IFMIF is conceived for fusion materials testing.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

REVOLUTION PROJECT: PROGRESS IN THE EVOLUTION OF SOLEIL MOTION CONTROL MODEL*

S. Zhang†, F. Blache, D. Corruble, C. Kheffafa, YM. Abiven, SOLEIL, Gif-sur-Yvette, France
S. Minolli, NEXEYA SYSTEMS, La Couronne, France

Abstract

SOLEIL is a third generation synchrotron radiation source located near Paris in France. REVOLUTION (REconsider Various contrOLler for yoUr motion) is the motion controller upgrade project currently in progress at SOLEIL. It was initiated to maintain the facility operations by addressing the risk of hardware obsolescence in motion control but at the same time making room for complex applications requirements to face new high performance challenges. In order to achieve these considerations, SOLEIL's strategy move was to go from a single controller for all applications to two motion controllers. A first Controller GALIL DMC-4183 was chosen to succeed the previous version DMC-2182. Both controllers can be integrated in the existing architecture with little hardware and software adaptation enabling full compatibility with the existing architecture. A second controller, Delta Tau Power Brick, has been selected as a HIGH PERFORMANCE solution providing advanced functionality.

The CLASSIC controller upgrade is about to be completed and the integration of Power Brick into the SOLEIL control system is ongoing. The system complexity is abstracted by embedding processing functions into low-level code and giving end-users a simple high-level interface. The work done to structure the interfacing and standardization of the controller are detailed in this paper.

CONTEXT OF REVOLUTION AT SOLEIL [1]

Motion Control Status and Upgrading Motivations

SOLEIL needs to upgrade its standardized GALIL motion controller in order to address two challenges: First, the current motion controller used in SOLEIL is dated and in risk of being discontinued and we must have an updated product before this happens. Nevertheless, this is not expected to occur in the next few years. Secondly, new motion control applications demand a higher performance and more advanced features that can not be reached by our current controller.

Strategy: Change Model

To achieve the next three goals of "Increasing performance, maintaining operational continuity, and controlling overall cost", SOLEIL decided to replace the current model of a "single and universal controller" by a model with two kinds of motion controllers:

- CLASSIC motion controller, mainly used and suitable for "simple and classic" applications. This controller is fully compatible with the existing hardware and software motion control architecture.
- HIGH PERFORMANCE controller, used in few cases, suitable for "complex and fast" applications. The hardware and software motion control architecture is designed to be close to the existing architecture for the CLASSIC controller.

CLASSIC CONTROLLER EVOLUTION

The GALIL DMC-4183 controller was chosen to succeed the current Galil DMC-2182. Its implementation needs some hardware and software developments which are currently being done to make it compatible with current architecture. On the software aspect: a new firmware was developed by GALIL which includes the new "Continuous Closed-Loop on Stepper motors" feature. It was validated by SOLEIL by an operational application done on the monochromaters of LUCIA and TEMPO beamlines. To use this feature, a new microcode (embedded code) was developed and validation tests are ongoing. For the hardware: in order to use the current ControlBox rack (same size, same pinouts), an interface board (MIG-4121) has been developed which adapt the DMC-4183 to the internal pinout of the ControlBox rack. The integration in the rack has been validated and is now operational.

HIGH PERFORMANCE CONTROLLER STANDARDIZATION

Controller Product

Delta Tau Power PMAC in the Power Brick LV-IMS format was selected for the new "HIGH PERFORMANCE" solution after a long evaluation process.



Figure 1: Power Brick LV-IMS.

Power PMAC is a general-purpose embedded computer with a built-in motion and machine-control application which offers: powerful computing capacity, multi-axis synchronization, encoder processing, virtual control through

* Work also supported by XT. Tran, M. Cerato, G. Renaud, E. Fonda and SAMBA Beamline staff, C. Engblom, Delta Tau Ldt., IMO JEAMBRUN AUTOMATION, Observatory-Sciences Ldt...
† szhang@synchrotron-soleil.fr

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

IBEX - THE NEW EPICS BASED INSTRUMENT CONTROL SYSTEM AT THE ISIS PULSED NEUTRON AND MUON SOURCE

M. J. Clarke, F. A. Akeroyd, K. Baker, G. Howells, D. Keymer, K. J. Knowles,
C. Moreton-Smith, D. E. Oram, ISIS, STFC Rutherford Appleton Laboratory, Oxon, UK
M. Bell, I.A. Bush, R.F. Nelson, K. Ward, K. Woods, Tessella, Abingdon, Oxon, UK

Abstract

Instrument control at ISIS [1] is in the process of migrating from a mainly locally developed system to an EPICS [2] based system. The new control system, called IBEX, was initially used during commissioning of a new instrument prior to a long maintenance shutdown. This first usage has provided valuable feedback and significant progress has been made on enhancing the system during the facility maintenance period in preparation for the move onto production use. Areas that will be of particular interest to scientists in the future will be linking feedback from live data analysis with instrument control and also providing a simple and powerful scripting interface for facility users. In this paper we will cover the architecture and design of the new control system, our choices of technologies, how the system has evolved following initial use, and our plans for moving forward.

BACKGROUND

The ISIS pulsed neutron and muon source has been providing world-class science since the mid-1980s and currently has over thirty beamline instruments producing world-leading research. The software and hardware used to control these beamline instruments plays a significant part in the productivity of ISIS.

The current instrument control system, known as the SECI system, consists of LabVIEW [3] drivers for equipment control alongside in-house developed software for scripting, data collection and experiment management, etc. This control system has performed well for many years; however, the construction of new beamline instruments and the refurbishment of older instruments has led to significantly more complex experiments becoming possible and, as a result, are pushing beyond the boundaries for which the current control system was designed. With this in mind, and with the prospect of even more complex instruments being built in the future, it was decided to create a new control system based on EPICS.

HARDWARE ARCHITECTURE

The IBEX control system runs on a Windows 7 x64 Virtual Machine (VM) which is hosted on a rack-mounted Windows 2012 server. Each beamline instrument has its own independent server. Windows was chosen for a number of reasons, such as:

- The controls team have significant experience and expertise with running control systems on Windows
- ISIS has an IT support group that can assist with installation and maintenance
- Easier interfacing to existing site infrastructure (e.g. data archiving, system backup/replication)
- Can use existing LabVIEW drivers more easily
- Ability to do a phased migration, running either in parallel or a mixed IBEX/SECI system

The beamline instrument server and equipment are usually placed on the general ISIS network; however, network sensitive equipment can be placed on a private network. This private network is implemented using a Virtual Local Area Network (VLAN) which can be routed back to the main computer room. As the virtual machines are replicated to the computer room it is possible to rapidly restore and run an instrument from there in the event of beamline server failure.

Most beamline and sample environment equipment in use on ISIS instruments are controlled via serial (RS232) communication. For ease of access and flexibility MOXA [4] NPort serial device servers are used to expose the serial devices to the network.

SOFTWARE ARCHITECTURE

A system based on EPICS lends itself naturally to a client-server style architecture and this is how IBEX has been implemented. Figure 1 shows the basic flow of information relating to the experiment; the dashed line indicates the separation between the client and server side components. Where possible the system uses EPICS Channel Access (CA) for communication between the client and server side.

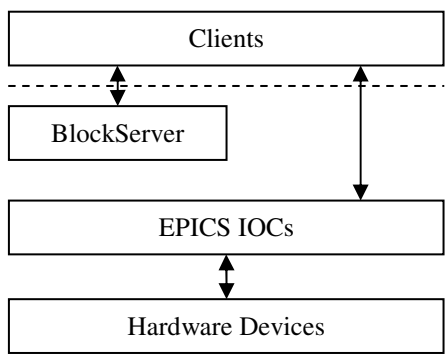


Figure 1: The flow of experiment information for a beamline instrument.

EPICS Input/Output Controllers (IOCs) are the standard way of communicating with the devices used on

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

100 Hz DATA ACQUISITION IN THE TANGO CONTROL SYSTEM AT THE MAX IV LINAC

P. Bell, V. Hardion, V. Michel, Max IV Laboratory, Lund, Sweden

Abstract

The MAX IV synchrotron radiation facility is currently being constructed in Lund, Sweden. A linear accelerator serves as the injector for the two storage rings and also as the source of short X ray pulses, in which mode it will operate with a 100 Hz repetition rate. The controls system, based on TANGO, is required to collect and archive data from several different types of hardware at up to this 100 Hz frequency. These data are used for example in offline beam diagnostics, for which they must be associated to a unique electron bunch number. To meet these requirements, the timing performance of the hardware components have been studied, and a TANGO Fast Archiver device created. The system is currently in the deployment phase and will play an important role in allowing the linac and Short Pulse Facility reach their 100 Hz design goal.

INTRODUCTION

The accelerator complex at the MAX IV laboratory consists of a 3 GeV, 250 m long full energy linac, two storage rings of 1.5 GeV and 3 GeV and a Short Pulse Facility (SPF). The repetition rate of the linac is a maximum of 10 Hz when serving as the injector for the rings and 100 Hz when providing pulses for the SPF.

The control system has a three-tier architecture, with specific hardware handling the real-time tasks and TANGO [1] representing the middle tier as the primary control system. Most equipment is interfaced directly to TANGO via TCP/IP. For the client layer, physicists and operators can interact with TANGO via its Python and MATLAB bindings or through the SARDANA [2] layer, which brings a macro server and standardised Graphical User Interfaces based on TAURUS [3].

In the operation of the linac it is necessary to be able to archive data through the TANGO system at the repetition rate or bunch frequency of the electrons. For example, oscilloscope devices connected to Current Transformers at various locations along the length of the linac are used to record the distribution in time of the charge of the passing electron bunches. These waveform data will be used, for example, in the offline diagnosis of any beam loss. The capturing of information from the control system at high frequency is referred to as Fast Archiving, since it must perform at up to 100 Hz to record every electron bunch when the linac is operating in SPF mode. The same system will also be used to read data from the linac and the storage rings at 10 Hz, where it will be used, for example, in the Slow Orbit Feedback system.

The work that is reported here has so far has focused on the integration of TANGO devices for controlling oscilloscopes into the Fast Archiver, though the inclusion of other equipment such as power supplies is foreseen in the future. The first section of the report documents how the performance of the many devices that will participate in the Fast Archiving is being systematically measured. The second section then describes the design and functioning of the Fast Archiver and how the oscilloscope devices have been adapted for inclusion in the system.

TIMING PERFORMANCE STUDIES

Timing performance studies are being performed on the TANGO devices that will take part in the Fast Archiving. The performance studies first check that the hardware itself is capable of sending or receiving commands at a sufficiently high frequency. An example is shown in Fig. 1 for a certain type of power supply used at Max IV.

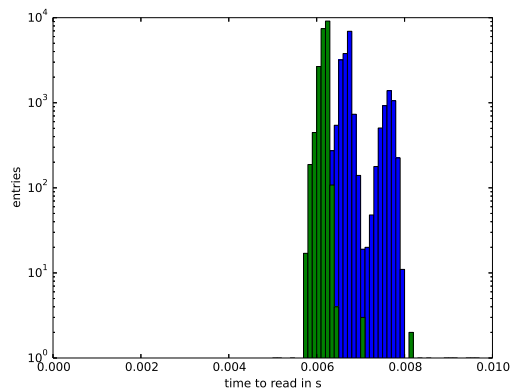


Figure 1: Distribution of time taken to read the current from a power supply over a direct raw socket connection (green) and via the "Current" attribute in its TANGO device integrated to the control system (blue).

A script is used that repeatedly reads, for example, the value of the current from the power supply hardware and plots the distribution of the time needed to receive a reply. In this example the distribution (in green) in well within the 10 ms envelope needed for 100 Hz operation. The performance of the TANGO device is then checked by measuring the distribution of the time needed to read the attribute "Current" once the device has been interfaced to the TANGO control system. The difference between the two distributions gives an indication of the overhead of the control system, including the network performance. In this case, the distribution (in blue) remains within the 10 ms envelope though adds a few ms to the "raw" performance of the hardware over a direct TCP/IP connection. It can be concluded that this hardware should be suitable for 100 Hz operation.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

TANGO-KEPLER INTEGRATION AT ELI-ALPS*

Péter Ács, Sándor Brockhauser, Lajos Jenő Fülöp, Veronika Hanyecz, Miklós Kiss, Csaba Koncz, Lajos Schrettner, ELI-ALPS, Szeged, Hungary

Abstract

ELI-ALPS will provide a wide range of attosecond pulses which will be used for performing experiments by international research groups. ELI-ALPS will use the TANGO Controls framework to build up the central control system and to integrate the autonomous subsystems regarding software monitoring and control. Beside a robust central and integrated control system, a flexible and dynamic high level environment could be beneficial. The envisioned users will come from diverse fields including chemistry, biology, physics or medicine. Most of the users will not have programming or scripting background. Meanwhile workflow system provides visual programming facilities where the logics can be drawn, which is understandable by the potential users. We have integrated TANGO into the Kepler workflow system because it gives a lot of actors for all natural scientific fields. Moreover it has the potential for running the workflows on HPC or GRID resources. We demonstrated the usability of the development with a beamline simulation. The TANGO-Kepler integration provides an easy-to-use environment for the users therefore it can facilitate e.g. the standardization of measurements protocols as well.

INTRODUCTION

ELI-ALPS is one of the three pillars of the European Extreme Light Infrastructure project. The Attosecond Light Pulse Source (ALPS) facility will provide a wide range of ultrafast light sources (such as coherent XUV and X-ray attosecond pulses) for performing material, condensed matter and surface science, chemical, biological, physical or medical experiments. Besides, the development of the technology for generating 200PW peak intensity pulses is also a main mission. As a research facility, the infrastructure will contain a large number of experimental devices and equipment which have to be managed and controlled by a robust and flexible system. We found that the TANGO Control system [1] is able to address this complexity, it has already been used at several large research infrastructures for more than a decade.

TANGO Controls is an open-source control system framework which provides a foundation to develop control systems. Architecture of a TANGO-based control system has three main components: the Device Server, the Database and the Client. The Device Server component is the base element that provides monitoring and control capabilities over a set of devices. The Database includes information about the Device Servers and Devices. A TANGO Client may connect to Device Servers using the reference acquired from the Database to perform actions on a certain Device. Experimentalist and scientist users may create programs by using the available APIs. However, not all of the users have a programming or scripting background necessary for such a task.

Scientific workflow systems [2-3] provide a graphical interface to create scientific applications without a deep scripting and/or programming background. A workflow is basically a collection of jobs connected to each other. Each job has input(s), performs a specific function and produces output(s) that can be used by other jobs. The jobs are usually represented as boxes with input and output ports. This kind of graphical representation of unit jobs increases productivity and helps maintainability. It allows the user to save the workflows as well as to share them with others.

The motivation of the paper is to develop a prototype integration of the Kepler scientific workflow system [4] with the TANGO framework. This way a high-level graphical environment could be provided for the users to create programs controlling and monitoring research equipment. Some other institutes already follow a similar approach [5-6]. However, this work has some further motivations, e.g. to collect relevant experience about how to integrate the TANGO framework into other environments.

THE INTEGRATION

A Kepler workflow has an arbitrary number of actors and exactly one director. Every actor represents a job with its necessary input and output ports, and an additional configuration dialog may help to fine-tune the behaviour. Actors communicate by sending messages via interconnected ports. The role of the director is to supervise the execution of the workflow, e.g. it manages the number of executions, enables parallelization if possible, etc.

A fundamental component of the TANGO Control System is the TANGO Device Server, which represents a set of devices with all of their attributes, commands, events and other relevant information. The goal was to create a universal *TANGO actor* in Kepler, which can dynamically adopt any functionalities (commands, attributes, etc.) a TANGO Device may have. However, it is implemented in a way that when the TANGO actor is instantiated, then the necessary functions can be chosen and will only be available afterwards.

It is not completely obvious how to represent and execute a TANGO workflow in Kepler because the functional structure of Kepler is fundamentally different from that of TANGO. As it was mentioned before, a Kepler workflow is controlled by directors.

* The ELI-ALPS project (GOP-1.1.1-12/B-2012-000, GINOP-2.3.6-15-2015-00001) is supported by the European Union and co-financed by the European Regional Development Fund.

† https://www.eli-alps.hu/?q=en

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

ELI-ALPS CONTROL SYSTEM STATUS REPORT *

Lajos Jenő Fülöp, Sándor Brockhauser, Sándor Farkas, Veronika Hanyecz, Miklós Kiss, Miklós Tamás Koncz, Árpád Mohácsi, Kwinten Nelissen, Lajos Schrettner, Balázs Szalai, Péter Szász, Chris Turner, ELI-ALPS, Szeged, Hungary

Abstract

ELI-ALPS will provide a wide range of attosecond pulses which will be used for performing chemical, biological, physical or medical experiments by international research groups. It is one pillar of the first international laser facility for the scientific user communities. ELI-ALPS use the TANGO Controls framework to build up the central control system and to integrate the autonomous subsystems regarding monitoring and control. It will be also used for the implementation of some autonomous systems' control system while others will be implemented differently. The central control system and the integration strategy of the autonomous systems is designed. The centralization and integration needs are surveyed and the requirements are collected. Prototypes have been developed to clarify the requirements and to test the designs. Requirements elicitation, designing and prototype development follows a Lean-Agile approach and includes several fields: device drivers and simulators; integration logic; central supervision, archiving, logging and error recovery; graphical user interfaces and so on.

INTRODUCTION

ELI-ALPS is one of the three pillars of the European Extreme Light Infrastructure project. The Attosecond Light Pulse Source (ALPS) facility will provide a wide range of ultrafast light sources (such as coherent XUV and X-ray attosecond pulses) for performing material, condensed matter, surface science, chemical, biological, physical or medical experiments. Besides, the development of the technology for generating 200PW peak intensity pulses is also a main mission. As a research facility, the infrastructure will contain a large number of experimental devices and equipment which have to be managed and controlled by a robust and flexible system. We found that the TANGO Control system [1] is a good candidate, it is already used by several research institutions, mostly in synchrotrons, but also in laser projects [2]. The basic layout of ELI-ALPS is the following: different primary laser sources generate laser pulses with different characteristics that are delivered by the beam transport lines to the secondary sources, which then generate attosecond light pulses for the experiments.

In the first phase of the development of the ELI-ALPS Control System the fundamental concepts, the high-level architectural design and the frames of the development process have been identified. In this paper the most important requirements, concepts, aspects, and prototypes are summarized briefly.

REQUIREMENTS

The general requirements are to follow public standards [3] as well as in-house policies and handbooks that also declare the application of frameworks such as TANGO and Taurus [4]. In order to maximize efficiency, different techniques are planned to be used. Automation techniques will be applied as much as possible, e.g. automatically registering devices [5], generate and configure GUIs and so on. Simulation environments, mimicking the devices and some key relationships will allow to perform most of the software development and testing independently from the real hardware.

Scientific Systems

Laser sources will be delivered as black-box, turn-key systems together with hardware and software. The secondary source systems are addressed by dedicated projects. The requirements, the technical design, as well as the hardware shopping list are provided by expert institutes of the corresponding secondary source field. The beam transport systems are addressed by a dedicated in-house project, which provides the requirements, the technical design and the hardware. Software and integration will be covered by another dedicated project.

Central Control System

The Central Control System will include the following services. The *archiving* subsystem collects all the preconfigured local or central variables for later usage and trends. The *alarms* subsystem collects and shows all the local and central alarms that are the result of evaluating predefined formulas on attribute values. The *logging* subsystem collects all the logs.

The *Integration Platform* supports, manages and orchestrates the collaboration of the systems (lasers, secondary sources, beam transport, etc.). Each system provides a gateway for communication. The gateways should be accessible only from the central system. For the other systems the central control system defines a uniform interface in order to provide controlled access to the gateways of a connecting system. A locking mechanism of the central system is responsible for system (proxy/gateway) allocation (i.e. granting exclusive access, when a system can be allocated to only one other system at a time). Moreover, it is preconfigured which systems may allocate a certain system for monitoring and control purposes.

* The ELI-ALPS project (GOP-1.1.1-12/B-2012-000, GINOP-2.3.6-15-2015-00001) is supported by the European Union and co-financed by the European Regional Development Fund.

† https://www.eli-alps.hu/?q=en

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

A FRAMEWORK FOR HARDWARE INTEGRATION IN THE LHCB EXPERIMENT CONTROL SYSTEM

L. Granado Cardoso, C. Gaspar, R.Schwemmer, J. Barbosa, F. Alessio, CERN, Geneva, Switzerland
P-Y. Duval, CCPM, Marseille, France

Abstract

LHCb is one of the 4 experiments at the LHC accelerator at CERN. During the upgrade phase of the experiment (planned for 2018), several new electronic boards and Front End chips that perform the data acquisition for the experiment will be added by the different sub-detectors. These devices will need to be integrated in the Experiment Control System (ECS) that drives LHCb. Typically, they are controlled via a server running on a PC which allows the communication between the hardware and the experiment's SCADA (WinCC OA). A set of tools was developed that provide an easy integration of the control and monitoring of the devices in the ECS. The fwHw is a tool that allows the abstraction of the device's models into the ECS. Using XML files describing the structure and registers of the devices it creates the necessary model of the hardware as a data structure in the SCADA. It allows then the control and monitoring of the defined registers using their name, without the need to know the details of the hardware behind. The fwHw tool also provides the facility of defining and applying recipes - named sets of configurations which can be used to easily configure the hardware according to specific needs.

INTRODUCTION

The LHCb experiment is one of the detectors collecting data at the LHC accelerator at CERN. It is specialized in b-physics and is composed of several sub-detectors and subsystems. All of these subsystems have specialized custom electronic boards and devices to acquire the data from the events produced by the collisions of the LHC beams in the detector. In order to easily and reliably control and configure these devices they need to be integrated in the SCADA System (WinCC OA) [1] so they can be driven by the Experiment Control System (ECS). Due to the custom nature of these devices, their description varies greatly between subsystems, so, to effectively integrate these devices in WinCC OA there was the need to create a tool that could abstract the description of the hardware, model it into the data structures used by WinCC OA and provide a layer to interface with this hardware.

THE FWHW TOOL

All the LHC experiments at CERN use as their SCADA system the software WinCC OA. In order to reduce duplication of development, the JCOP (Joint Controls Project) [2] was created in order to provide common controls solutions for the four experiments. JCOP provides a framework for the creation of JCOP components – WinCC OA packages containing all the required user panels, libraries, scripts and other software – that can be easily installed and distributed.

The framework Hardware tool (fwHw) [3] is developed as a JCOP component, allowing for easy installation by the sub-detectors. The tool provides a user interface, which allows for the easy modelling and configuration of the existing hardware devices and can also serve as a debugging tool. However, its purpose is mostly setting up the system and the different sub-detectors or subsystems should then implement their hierarchical control Finite State Machine (FSM) trees and their own user interfaces more suitably geared to the operation of their electronics.

Hardware Abstraction

The fwHw tool models the hardware into WinCC OA datapoint structures. The WinCC OA data structure is a tree-like structure, where a logical model is defined in Datapoint Types, which can be instantiated in datapoints that share this logical model.

Similarly the fwHw tool produces the logical model of the electronic devices as Datapoint Types, but providing an interface more adapted to the modelling of electronic devices.

In the fwHw tool, the data is hierarchical organized with the following types:

- Registers.

The registers are the representation of a register on a real hardware device. A register can be of any size and implement any type of protocol. The hardware tool allows the configuration of each register according to the settings necessary to access the hardware via the different protocols, the size of the register and the type of readout it will require.

The registers can also represent a local variable, i.e. a logical data item of a given area (e.g. a string which holds a filename to configure a given area).

These are the leaves or the lower level nodes of a hardware device model tree.

- Areas/Sub-Areas

Areas are logical groups of registers and/or other areas, which represent a logical division of a part of the hardware device (e.g. a chip or a group of chips). Defined areas can be declared as sub-areas of other areas and can be used multiple times. For example, in Figure 1 we can see that both areas 2 and 3 have the same type and thus have the same structure. The defined areas can also be used by different Device Types, so if some devices share a common structure, there is no need to duplicate similar structures.

- Device Types

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

SYNCHRONISING HIGH-SPEED TRIGGERED IMAGE AND META DATA ACQUISITION FOR BEAMLINES

N. De Maio, A. P. Bark, T. M. Cobb, J. A. Thompson
Diamond Light Source Ltd., Didcot OX11 0DE, UK

Abstract

High-speed image acquisition is becoming more and more common on beamlines. As experiments increase in complexity, the need to record parameters related to the environment at the same time increases with them. As a result, conventional systems for combining experimental meta data and images often struggle to deliver at a speed and precision that would be desirable for the experiment. We describe an integrated solution that addresses those needs, overcoming the performance limitations of PV monitoring by combining hardware triggering of an ADC card, coordination of signals in a Zebra box [1] and three instances of areaDetector streaming to HDF5 data. This solution is expected to be appropriate for frame rates ranging from 30Hz to 1000Hz, with the limiting factor being the maximum speed of the camera. Conceptually, the individual data streams are arranged in pipelines controlled by a master Zebra box, expecting start/stop signals on one end and producing the data collections at the other. This design ensures efficiency on the acquisition side while allowing easy interaction with higher-level applications on the other.

INTRODUCTION

The kinds of data rates coming out of a typical modern synchrotron beamline continue to push up against the current limits in experimental complexity, processing power and storage space. Ongoing improvements in automation and computing infrastructure help scientists get past those limits, decreasing the time individual scans of a sample take to set up and carry out.

When it comes to taking data, a number of well-established technologies and beamline designs provide stable frameworks to acquire as much as possible across a number of different devices and channels. Each of these comes with its own requirements with regards to configuration and operation. Each of these may or may not also provide its own native mechanisms to sample data at user-defined regular intervals.

Getting these devices to perform their measurements in a synchronised way poses some unique challenges. If the samples are taken a sufficient time apart, the middleware layer provided by the controls software can, and often does, take on the role of coordinator. But there is a point where the interval between two samples gets too small to guarantee accurate time stamps. Different devices may also expect different kinds of triggers, and every additional layer of software mediating between those makes the delays between the data acquisition system and the hardware more unpredictable. Figure 1 illustrates this problem.

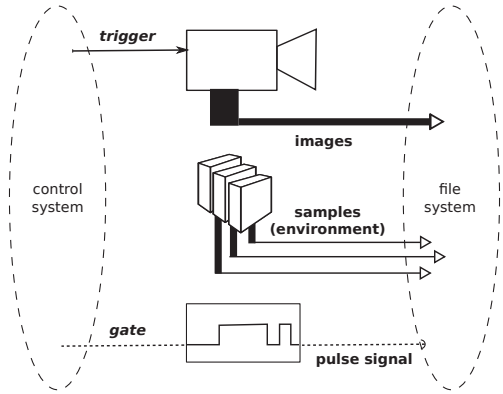


Figure 1: Unsynchronised trigger and data flow during image acquisition.

ARCHITECTURE

The answer to coordinated sample acquisition at small time scales is to use hardware triggering. Using a Zebra box [1], we present a design for a fast data acquisition system intended for sampling rates within the range of 30Hz and 1kHz. It not only coordinates the different hardware triggers involved but also bundles the acquired samples in three parallel pipelines, resulting in three files containing measurements for the same time frame.

Figure 2 gives an overview of the architecture. As Diamond Light Source (DLS) uses EPICS [2] as its control system, we rely on the areaDetector framework [3] and the EPICS implementation of the Zebra driver for configuring and collecting data from the underlying devices. We also rely on areaDetector plugins to post-process this data to some extent and to redirect it to files.

The control signal flow in our model is driven by the Zebra box. In our beamline-specific application, it receives an external trigger signal of its own. That signal causes the Zebra to produce one or more outgoing pulses, which are then forwarded to the connected devices. The incoming pulse need not be another hardware signal. Arm or Acquire commands to the relevant PVs are just as suitable because it is the configuration of the FPGA logic in the Zebra that ultimately decides what kinds of pulses to send out in response, and how many.

If individual devices expect different pulse shapes as trigger signals, the Zebra box acts as a mediator between incoming and outgoing triggers while ensuring that all outgoing pulses are sent off at the same time. If they are capable of receiving external triggers, they should be set to operate in this way. If they are not, the trigger signal needs to be recorded on one of the data channels so it can be extracted

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

QUICK EXPERIMENT AUTOMATION MADE POSSIBLE USING FPGA IN LNLS

M. P. Donadio*, H. D. Almeida, J. R. Piton
CNPEM / LNLS, Caixa Postal 6192, Campinas - 13083-970, Brazil

Abstract

Beamlines at LNLS are being modernized to use the synchrotron light as efficiently as possible. As the photon flux increases, experiment speed constraints become more visible to the user. Experiment control has been done by ordinary computers, under a conventional operating system, running high-level software written in most common programming languages. This architecture presents some time issues as computer is subject to interruptions from input devices like mouse, keyboard or network. The programs quickly became the bottleneck of the experiment. To improve experiment control and automation speed, we transferred software algorithms to a FPGA device. FPGAs are semiconductor devices based around a matrix of logic blocks reconfigurable by software. The results of using a NI Compact RIO device with FPGA programmed through LabVIEW for adopting this technology and future improvements are briefly shown in this paper.

INTRODUCTION

Experiments at LNLS historically were made using software running in conventional operating systems. Our first software for beamline control centralized all the operations, sensor reading and actuators controlling in a single thread. This was the cause of lots of disturbances in motor movement, when the user moved a mouse or selected a menu item, which almost never ran in a smooth pattern. We could see also a big dead time caused by many tasks running sequentially rather than in parallel.

In 2010 we started to use EPICS [1] and, this way, we could decentralize all the access to the devices. Doing so, we could get many improvements in motor controlling and reduction of dead time. The system became easier to maintain as each device is accessed by a single piece of stan-dalone software. Another advantage was to provide the infrastructure to build a web system named LabWeb [2], based on Science Studio [3] developed and applied at the Canadian Light Source (CLS), that became possible for users to operate the beamline from their universities.

From 2010 until recently in the current year not only the software, but most of the control hardware was replaced in all beamlines. After some years developing, improving and using the new EPICS based system, we faced a new challenge: the beamlines had its photon fluxes and detector sensitiveness increased by the hardware improvement and software was again the bottleneck for some experiments. EPICS can provide a good experiment control using only software and not firmware, in conventional operating systems, if there is no concern about times lesser than 100 ms. Control loops quicker than 10 Hz are not guaranteed to be attended inside a good error limit. To go beyond this limit we needed to change the paradigm adopting a real-time system, for example, or implementing the fast piece of the experiment in hardware.

The first beamline to see this limitation was IMX [4], followed by XRF [5]. Both needed to synchronize motor position, detector shot and shutter in a few milliseconds with error margin of some microseconds.

The third beamline case was SAXS1 [6], whose system is presented in this paper. The development of the system for SAXS1 was used also to test the approach to the beamlines that will be built in Sirius [7], the new Brazilian Synchrotron. The timing for Sirius beamlines needs to be in the microsecond loop or faster.

SAXS1 HARDWARE

We need a system with low jitter between the trigger signal and the start and the end of data acquisition. Going near the hardware is an approach to reduce jitter and that is the motivation to choose FPGA technology. We had available in our inventory a cRIO NI 9144 [8], a unit with FPGA, so this was a natural choice to use in the work. cRIO NI 9144 can be connected to a PXI [9] only by an EtherCAT connection.

Figure 1 shows the connection between the cRIO and detectors and actuators.

There are two main detectors, a Pilatus [10] 300K to measure SAXS and a Pilatus 100K to measure WAXS. Pilatus 300K is configured to collect data in a determined frequency and sends a trigger TTL signal every time an acquisition is being done. cRIO is responsible to send a TTL trigger to Pilatus 100K as it can start acquisition.

Two photo-diodes provide the beam intensity data through a Stanford SR570 [11]. One of the photo-diodes is located in the beam stopper and receives radiation after the sample. The other one is located before the sample and receives the reflection of a small part of the beam thanks to a mylar or kapton thin film positioned in a 45-degree angle related to the beam. cRIO reads the voltage data from Stanford SR570 using a NI 9215 module [12].

cRIO also controls a fast shutter to limit to a minimum the time the sample is irradiated, by opening the shutter only when Pilatus is acquiring.

A cRIO is used to:

- run the FPGA code that sequences the experiment
- read voltage from Stanford using a NI 9215 module

* marcio.donadio@lnls.br

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

NEUTRON SCATTERING INSTRUMENT CONTROL SYSTEM MODERNIZATION - FRONT-END HARDWARE AND SOFTWARE ADAPTION PROBLEMS

M. Drochner, L. Fleischhauer-Fuss, H. Kleines, M. Wagener, S. v. Waasen
FZ Jülich / ZEA-2, Jülich, Germany

Abstract

When the FRM-2 neutron source went into operation (2002) and many instruments were moved from the closed-down Juelich reactor to the new facility, it was agreed on a choice of front-end hardware and the TACO middleware from ESRF. To keep up with software standards, it was decided recently to switch to TACO's successor - the TANGO control software. For a unified "user experience", new graphical user interface software "NICOS-2" is being developed by the software group at FRM2.

While general semantics of TACO and TANGO don't look very different at a first glance, and adaption of device servers seemed to be straightforward at first, various problems in practical operation were found, due to difference in state handling, timing behavior and error reporting. These problems, and the changes that had to be made to ensure reliable operation again, will be described.

INTRODUCTION

After some different developments turned out to be rather short-lived, partly due to political reasons, partly due to manpower issues, the "NICOS2" instrument control and user interface framework got administrative backing and developer personnel to be further maintained and adopted by neutron scattering instruments at FRM-2. (See [1] and [2] for some historic information.)

For device access, it was decided to abandon TACO and rewrite device servers to use the more modern TANGO framework. Since the basic functionality of TANGO device servers, the remote, synchronous, execution of commands, is quite similar to that of TACO, an easy migration was expected. It was assumed more or less that the old implementation of device server commands just needed to be adapted to a slightly different API syntax.

BOUNDARY CONDITIONS, DEVICE SERVER ARCHITECTURE

A number of device servers, in particular those which handle detectors, have a structure as shown in Fig. 1.

It is often a matter of discussion whether detector data which come as individual events should be histogrammed immediately during data acquisition. It is not subject of this paper to discuss this design decision – it should just be considered that real-time histogramming needs to be done anyway, because instrument operators want a live picture of measured data. If single event data provide no additional advantage, the cost (both in terms of runtime overhead and storage media) can be saved. Up to now, no compelling reason for single event data was found.

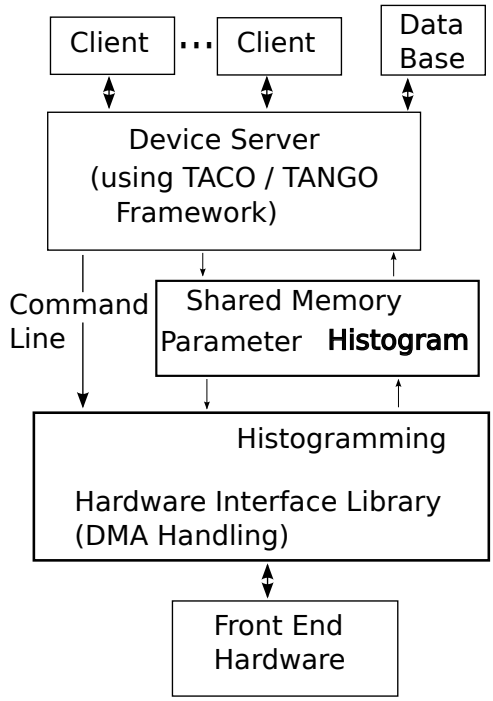


Figure 1: Structure of a detector device server which does real-time histogramming of incoming data.

Since TACO is not multithreaded from the beginning, and very likely not prepared to run with multiple threads active (actually, it was seen some years ago that even the underlying RPC library was not thread-safe), detector access was put into a separate process which runs asynchronously to TACO command execution. Another advantage of this was that the detector code could be run from a standalone process, for tests when the TACO manager and database services were not available.

Communication between the server core and the detector access program uses shared memory — keeping a table telling detector parameters (e.g. TOF slot settings) in one direction and the detector data histogram in the other. The shared memory was implemented as a memory-mapped file. This was helpful for debugging because the histogram data could be looked at everytime just using standard UNIX command line tools like od(1)/hexdump(1).

NICOS is a comfortable user interface to control the instrument, define and run measurement scans, show a life display of detector data and support logging and event han-

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

THE NEW TANGO-BASED CONTROL AND DATA ACQUISITION SYSTEM OF THE NEUTRON SPECTROMETER DNS AT FRM II

H. Kleines, M.Drochner, M.Wagener, L. Fleischhauer-Fuss, S. Keuler, F. Suxdorf, R. Möller, S. Janasche, S. van Waasen, K.-H. Mertens, M. Bednarek, K. Bussmann, Y. Su,
Forschungszentrum Jülich, 52425 Jülich, Germany

Abstract

Forschungszentrum Jülich has been operating the neutron instrument DNS at the neutron source FRM II for about 10 years. DNS is a time of flight neutron spectrometer with polarization analysis that experienced a major upgrade in 2014 and 2015. During the upgrade DNS was equipped with new electronics and a new control and data acquisition system, including a transition from the existing TACO system to its successor TANGO. On the client side the NICOS software developed at FRM II is used for the implementation of measurement operations and user interface. The design of the new control and data acquisition system is presented and the lessons learned by the introduction of TANGO are reported.

INTRODUCTION

Forschungszentrum Jülich developed and operates about 20 neutron instruments at its outstations ILL in Grenoble, the Spallation Neutron Source in Oak Ridge and at the FRM-II in Garching near Munich. The control and data acquisition systems of almost all these neutron instruments are based on the so-called "Jülich-Munich Standard", a joint effort of ZEA-2 (central electronics institute of Forschungszentrum Jülich) and Technical University Munich (TUM) to define a common framework for the electronics and software of neutron instruments that is followed by most instruments at the FRM-II [1]. It is based on the TACO [2] control system developed by the ESRF and the extensive use of industrial type front-end equipment, e.g. PLCs, fieldbus systems or remote I/Os. Because TACO is considered to be outdated now, there was the decision to introduce its successor TANGO [3] at FRM-II as joint effort between TUM and ZEA-2. Since the neutron instrument DNS experienced a major upgrade it was selected as one of the pioneer instruments for the introduction of TANGO, simultaneously with the neutron instrument BIODIFF. At the same time the new instrument control software NICOS (developed by TUM) was implemented on both instruments.

THE "JÜLICH-MUNICH STANDARD"

The "Jülich-Munich standard" is a framework for the selection of technologies and components at each level of the control system. The definition of this framework was motivated by synergy effects and the reduction of spare parts on the shelf. A guiding principle for the framework was to minimize the development efforts and to acquire as much from the market as possible. A key component of the framework is the consistent use of industrial technologies like PLCs, fieldbus systems or decentral periphery in the front end. Main motivations are:

- low prices induced by mass market
- inherent robustness
- long term availability and support from manufacturer
- powerful development tools

A control system according to the Jülich-Munich Standard is organized hierarchically into the following levels:

Field level: The field level is the lowest level, at which devices that are not freely programmable reside, like motor controllers, SSI controllers, PID controllers, analogue and digital I/O modules, or measurement equipment. For all industrial type of I/O modules PROFIBUS DP based decentral periphery is recommended. Siemens ET200S is the preferred one. The ET200S modules 1STEP and 1STEPdrive are the predominantly used the stepper motor controllers.

Control level: The control level resides on top of the process level. Devices at the control level are freely programmable. They must meet real time requirements and guarantee robust operation in a harsh environment. At the control level Siemens S7 PLCs are used, because they dominate the European market.

Process communication: Process communication covers the communication of devices at the field and control level with supervisory controllers or computers. For lab equipment GPIB and proprietary RS232/RS485 connections are unavoidable. For industrial automation equipment PROFIBUS DP and PROFINET are the recommended choices. They are the dominating fieldbus systems in Europe and naturally supported by S7 PLCs and many other devices. A major reason for their success is the technological and functional scalability based on a common core as well as the programming model, which easily maps to PLC operation.

Experiment Computer: For economic reasons, all experiment computers should be PCs. Linux, being well established in the scientific community, is the only supported operating system. There is no definition of a specific kernel version or distribution. Direct device access should not be implemented on conventional PCs but on CompactPCI systems. CompactPCI allows deploying a variety of existing software in a mechanically more robust platform that fits into 19" racks.

Middleware: Since the framework aims at an inherently distributed system, software support for the transparent distribution of services between systems is

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

MOTION CONTROL ON THE MAX IV SOFT X-RAY BEAMLINES WITH TANGO AND SARDANA

M. Lindberg, J. Forsberg, L. Kjellsson, A. Milán,
C. Såthe, P. Sjöblom, S. Urpelainen, MAX IV, Lund, Sweden

Abstract

MAX IV Laboratory, a synchrotron facility in Lund, has selected TANGO as the control system framework for the entire facility. On the beamlines that are being built the Python-based SCADA (supervisory control and data acquisition) system Sardana will be used for experimental control.

SPECIES, one out of eight new soft X-ray beamlines, is used as a test bench for evaluating the chosen standards. Sardana is used to control the energy setting of the PGM (plane grating monochromator) as well as to provide macros and other utilities for the user. Generic Taurus GUIs and a SVG-synoptic give the user a way to interact with the control system and display relevant information. The standardized graphical interfaces give a familiar look and feel across the entire facility.

All motorized axes are controlled with the IcePAP motion controller. For the axes of the PGM, the IcePAP driver operates in hardware closed loop. Special care is taken in order to avoid slow and inaccurate movements of the PGM energy due to the non-linear relationship between the motors and the angular encoders.

INTRODUCTION

MAX IV Laboratory [1] will be a cutting edge synchrotron located in Lund, Sweden, and it is currently commissioning its main storage ring. This new facility builds on the experience from the previous MAX-lab in Lund, which is being decommissioned after 25 years of successful service to users from all over the world.

At this stage fourteen beamlines are funded and some are already in the process of being assembled at the MAX IV. The first users are expected in 2016. The SPECIES soft x-ray beamline [2, 3] is currently commissioned and taking users at the older facility MAX-lab with the intention to serve as a test bench for the control system that will be used at all the other beamlines at MAX IV.

During the commissioning of the SPECIES beamline, several motorized optical elements needed to be controlled to adjust the alignment of the beamline and its two branches, 40 meter long in total, as seen in Fig. 1. Also, the characteristics of the photon beam must be well known. This is achieved by using a set of macros making calculations from motor positions and current measurements.

MAX IV BEAMLINE CONTROL SYSTEM

By making use of control system standards already in use at other facilities, it is ensured that there are several tools such as pseudomotors, predefined scans and generic GUIs available from the start of the beamline commissioning. For the majority of the motorized axes the motion depends linearly on the corresponding encoder which makes it straight forward to achieve precise closed loop control but the PGM requires more effort which is further described below.

Control System Standards

The control system for the MAX IV facility [4] applies several standards for both accelerators and beamlines, that have been adopted with the aim to make the deployment and maintenance of a complex machine and a large number of beamlines more efficient and also to grant some familiarity to the user community.

Some of these standards are the TANGO Controls toolkit [5], the beamline SCADA Sardana [6], and the PyQt based Taurus GUI framework [7]. The IcePAP motion controller developed at the ESRF [8] has been chosen as a facility wide standard. Other common references and tools have been developed at MAX IV such as the MAX IV unified coordinate system and the SVG-synoptic interface [9].

Tango, Sardana and Taurus are all open source software and, since they have been developed and used at several other synchrotron facilities such as ESRF, Alba, and Soleil, code and experience have been accumulated over time. To deepen the transfer of knowledge and to contribute to the community, MAX IV is a part of the Tango collaboration, the Sardana collaboration, and the IcePAP collaboration [10].

Commissioning with the Control System

During the development a generic Taurus GUI was used to access motors and configure them, to run predefined scans and modify the scan parameters. The generic Taurus GUI is fast to set up and flexible to use. During the progression of the beamline installation and commissioning new panels for each equipment were added as each set of motors was configured. An example can be seen in Fig. 2.

After the commissioning a synoptic like the one in Fig. 3 will replace the development GUI. The synoptic gives a good overview of the beamline based on the physical location of the equipment. It also shows alarm states and allows quick access to detailed panels for all equipment when the user clicks the corresponding icon. Commonly used functions such as acquisition GUIs may also be launched from the synoptic.

MOTION CONTROL

Motion axes at MAX IV, including the ca 60 axes of SPECIES, will mostly be motorized using 2-phase stepper motors controlled by the IcePAP motion controller [11]. Sardana interfaces well with IcePAP and provides the possibility

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

SYNCHRONIZED RAMPING OF MAGNET POWER SUPPLIES FOR STREAMLINED OPERATION AT ENERGY RECOVERY LINAC (ERL) AND ELECTRON LENS (e-Lens)

P. Kankiya†, J. Jamilkowski, T. Samms
Brookhaven National Laboratory, Upton, NY 11973 USA

Abstract

Synchronous ramping of an assembly of magnets is critical for operation of beam in an accelerator. Magnet currents must remain within the operational limits to avoid dis-alignment of electron beam. In order to comply with the design specifications of ERL and e-Lens project, two different software control mechanisms have been developed. The ramp profile is automated and maintained by tracking current in all dipole magnets at ERL and superconducting solenoid magnets at e-Lens. This mechanism speeds up operations and adds a level of protection. The purpose of this application is to reduce unnecessary interlocks of the personnel protection system. This paper will describe the power supply arrangement, communication mechanism and the state machine algorithm used for feedback and control. A report on operating experience will be presented.

INTRODUCTION

In year 2015 Energy Recovery LINAC project went under commissioning and their operations required design of special magnet ramp sequence to ensure compliance with accelerator design. The Energy Recovery LINAC (ERL) is designed to serve as a test bed for future projects such as LeReC, eRHIC and CeC. The lattice is unique and has been designed with a significant amount of flexibility in order to cover the required operational parameter space [1]. Another accelerator project with special ramp profile requirements is the RHIC electron lens (e-Lens) [2]. The Yellow and Blue e-Lens assemblies have a collection of 17 superconducting magnets. Current reference waveforms for these solenoids are generated by high level software applications and are preloaded to the VME based Waveform Function Generator (WFG) modules which are custom boards developed in-house.

ERL MAGNET MANAGER

ERL's R&D prototype is currently installed with an assortment of magnets including several quadrupoles, dipoles and solenoids shown in Fig 1. The large beam size and very low longitudinal emittance of the e-beams dictate that the tolerances on the magnetic fields are very tight. Significant misalignment in the electron beam could lead to stray radiation; therefore very strict radiation safe functional bounds are enforced by an independent system called the Particle Accelerator Safety System (PASS). This system is designed to protect personnel working in the areas near ERL by interlocking critical devices under specific abnormal conditions. In order to avoid unnecessary PASS interlocks due to mismatched critical dipole magnet currents, it was advised by the external reviewers of the project that software controls be modified. The intent is to reduce accelerator downtime, since each interlock of this type requires a sweep of the blockhouse by trained personnel. Another role of this software that has been added to the operational requirements is to manage magnet power supply controls and ERL e-beam orbit feedback corrections.

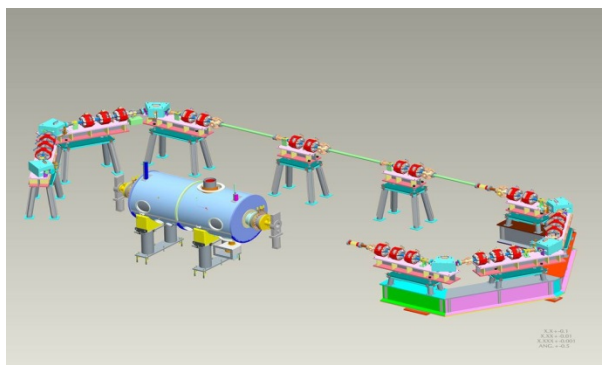


Figure 1: 3-D schematic of magnets in ERL beam line.

Software Implementation

Different COTS power supply systems with small amount of customisation are purchased and mounted to interface with various magnets based their on functional specifications. These supplies communicate over TCP/IP by means of custom multiple types of software developed using the Accelerator Device Object ADO framework [3].

"erlMagMan" is a program consisting of two levels of object oriented entities controlling power supplies which interface directly with the magnets. The first level is a an ADO class representing each magnet and consists of magnet properties such as bending angles of each dipole and nominal current at designed bending angle. These coefficients are provided by system experts and are obtained by direct measurements or simulations. Desired magnet currents are calculated at a given level of energy or at adjustment of trim angles. This energy will be entered manually by the shift operator or through the orbit correction program. Reverse conversion logic for translating magnet current to correction angles is also in place.

The second ADO layer acts as a governing class which interacts with all magnet ADOs simultaneously. It broadcasts the change of energy so all power supplies are configured with the correct current set point at the same time.

*† pkankiya@bnl.gov *Work supported by Brookhaven Science Associates, LLC under Contract No. DE-AC02-98CH10886 with the U.S. DOE

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

MeerKAT CONTROL AND MONITORING SYSTEM ARCHITECTURE

Neilen Marais*, SKA South Africa, Cape Town

Abstract

The 64-dish MeerKAT radio telescope, currently under construction, comprises several loosely coupled independent subsystems, requiring a higher level Control and Monitoring (CAM) system to operate as a coherent instrument. Many control-system architectures are bus-like, clients directly receiving monitoring points from Input/Output Controllers; instead a multi-layer architecture based on point-to-point Karoo Array Telescope Control Protocol (KATCP) connections is used for MeerKAT. Clients (e.g. operators or scientists) only communicate directly with the outer layer of the telescope; only telescope interactions required for the given role are exposed to the user. The layers, interconnections, and how this architecture is used to meet telescope system requirements are described. Requirements include: Independently controllable telescope subsets; dynamically allocating telescope resources to individual users or observations, preventing the control of resources not allocated to them; commensal observations sharing resources; automatic detection of, and responses to, system-level alarm events; high level operator controls and health displays; automatic execution of scheduled observations.

INTRODUCTION

The MeerKAT Control and Monitoring (CAM) subsystem is the high level control system that enables all the other subsytems to operate cohesively as a single telescope. Most of the CAM design principles were presented at ICALEPCS 2013 [1]. While much detail design and implementation has been done since then, the only major architectural addition to CAM is the concept of telescope subarrays and the subarray-manager components that implement them. An unusual feature for the control architecture of a large scientific instrument is the lack of any messaging middleware or message bus.

A brief MeerKAT status update is given, and a high level overview of the MeerKAT CAM architecture is presented with a description of how it is used to meet CAM requirements in practice. The scalability and reliability of this architecture are discussed.

MeerKAT DESCRIPTION AND STATUS

MeerKAT will be a 64-receptor[1] aperture-synthesis inteferometric radio telescope array. Receptor antennas are offset Gregorian antennas with a 13.5m main reflector in a feed-low configuration. An offset optical configuration has been chosen because its unblocked aperture provides optimal optical performance and sensitivity with good rejection of unwanted radio frequency interference from satellites and terrestrial transmitters. It also enables the installation of multiple receiver systems in the primary and secondary focal areas and provides a number of operational advantages.

The antenna platform supports up to four receiver systems mounted on a carousel, allowing the telescope to switch frequency bands within minutes. The initial design committed to cryo-cooled L-band receivers. Some experimental uncooled K-band receivers are installed during commissioning. Since then, the MeerKAT project has committed to installing UHF-band receivers and the Max-Planck-Institute for Radio Astronomy has committed to funding and building S-band receivers for MeerKAT.

MeerKAT receptors directly digitise the received signal, without transmitting analogue signals to a central facility – a first for radio telescopes. A central Time and Frequency Reference (TFR) subsystem using a GPS-disciplined Rubidium clock provides a highly accurate absolute time reference and a sub-picosecond accurate clock signal for phase-coherent digitiser operation. The clock and timing pulses are distributed to the receptors via dedicated fibre links. The TFR will later be upgraded to a MASER timing source.

Signals are transmitted from from the receptor digitisers via multiple bonded 40 Gigabit Ethernet (40GbE) fibre links to the Correlator-Beamformer (CBF) subsystem in the Karoo Array Processor Building (KAPB) using the packetised SPEAD [2] format. A separate physical Ethernet network is used for CAM traffic from the receptors.

The CBF subsystem performs the first level of signal processing and data reduction using the SKARAB open FPGA platform [3], and multicast capable commodity 10/40GbE switches as a data fabric. It produces correlated visibilities for imaging, and beam-former voltage data that is consumed by the Science Processing (SP) subsystem and optionally by user supplied equipment provided by third parties. SP is a software subsystem running on general purpose computers in the KAPB, utilising GPU- and/or other forms of acceleration. SP processes the CBF output into images and other relevant science outputs, and also archives observation data.

Ancillary subsystems include: the MeerKAT site Camera subsystem; the weather monitoring subsystem; the KAPB building management subsystem that monitors ambient conditions and RFI-door intrusion detection where appropriate; the Power Distribution Units (PDUs) of several subsystems to facilitate emergency power shutdowns. These subsystems are not directly involved in observations but are monitored to ensure safe and reliable operation of the telescope.

All subsystems are coordinated and monitored by CAM using a logically separate CAM network, using Ethernet as a fieldbus. Communication with all other subsystems is via the Karoo Array Telescope communications protocol (KATCP) [4]. CAM Device Translators are used to provide a KATCP interface where this is not available natively.

* nmarais@ska.ac.za

[1] A receptor logically groups the subsystems of a single antenna dish.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

REPORT ON CONTROL/DAQ SOFTWARE DESIGN AND CURRENT STATE OF IMPLEMENTATION FOR THE PERCIVAL DETECTOR

A.S. Palaha, C. Angelsen, Q. Gu, J. Marchal, N. Rees, U.K. Pedersen, N. Tartoni, H. Yousef, Diamond Light Source, Harwell Science and Innovation Campus, Oxfordshire, UK
M. Bayer, J. Correa, P. Gnadt, P. Gottlicher, H. Graafsma, S. Lange, A. Marras, S. Reza *, I. Shevyakov, S. Smoljanin, J. Supra, M. Tennert, U. Trunk, C.B. Wunderer, Q. Xia, M. Zimmer, Deutsches Elektronen-Synchrotron (DESY), Hamburg, Germany
D. Das, N. Guerrini, B. Marsh, T. Nicholls, I. Sedgwick, R. Turchetta, Rutherford Appleton Laboratory (RAL), Harwell Science and Innovation Campus, Oxfordshire, UK
G. Cautero, D. Giuressi, A. Khromova, R. Menk, G. Pinaroli, L. Stebel, Elettra Sincrotrone Trieste, Italy
H.J. Hyun, K.S. Kim, S. Rah, Pohang Accelerator Laboratory, Pohang, Korea

Abstract

The Percival Collaboration is developing a high-speed, low X-ray energy detector capable of detecting single photons while maintaining a large dynamic range of sensitivity.

The increased brilliance of state-of-the-art Synchrotron radiation sources and Free Electron Lasers require imaging detectors capable of taking advantage of these light source facilities. The PERCIVAL ("Pixelated Energy Resolving CMOS Imager, Versatile and Large") detector is being developed in collaboration between DESY, Elettra Sincrotrone Trieste, Diamond Light Source and Pohang Accelerator Laboratory.

It is a CMOS detector targeting soft X-rays < 1 KeV, with a high resolution of up to 13 M pixels reading out at 120 Hz, producing a challenging data rate of 6 GiB/s.

The controls and data acquisition system will include a Software Development Kit to allow integration with third party control systems like Tango and DOOCS; an EPICS [1] areaDetector [2] driver will be included by default. It will make use of parallel readout to keep pace with the data rate, distributing the data over multiple nodes to create a single virtual dataset using the HDF5 file format for its speed advantages in high volumes of regular data.

This development project will culminate in a control and DAQ system capable of dealing with very high data rates while providing easy integration with site-specific control systems.

This report presents the design of the control system software for the Percival detector, an update of the current state of the implementation carried out by Diamond Light Source.

INTRODUCTION

The PERCIVAL detector is designed for high-speed and low noise detection of soft X-rays. The Percival CMOS chip is a large pixel array of ~13 mega-pixels with 15 bits per pixel (packed in 16 bit words) which encodes a substantially higher dynamic range of the sensor as described below [3]. This leads to 25 MiB per frame. For each data frame the detector will also produce a reset frame of the same dimension, so the total data produced is 50 MiB per frame. There will also be a ~2 mega-pixels version [4]. For the designed frame-rate of 120 Hz this leads to an overall data rate of ~6 GiB/s [5].

The Detector

The detector pixel is an adaptive gain design [4], allowing single photon discrimination at low flux while still being capable of high flux measurements at the cost of increased noise. It comprises a diode and three capacitors, each of increasing capacitance, to collect charge depending on whether the next smallest capacitance component has been filled; thus "adaptive-gain". As there are four "collectors" per pixel, there are four possible gain values also. A pixel array feeds scrambled data to a pair of "mezzanine" cards, each with an FPGA to perform a certain amount of low level unscrambling and formatting the data into UDP packets for transfer to the processing nodes. These nodes are a Linux cluster of 8 commodity servers, as recommended by an earlier feasibility study [3] in order to handle the processing of the expected data rate.

CONTROL SYSTEM ARCHITECTURE

Shown in Fig. 1 is an overview of the proposed Percival software architecture, where the green single direction arrows represent the flow of data and purple bidirectional arrows represent control links. The top half represents the data acquisition system, with the sensor board having its data collected by the carriers boards via LVDS lines and transferred to the mezzanine boards before it is sent through the deep buffer switch and onto the Linux cluster.

Each FPGA of the mezzanine board has four 10 GBps links to the deep buffer switch, one for control and three for transmitting data. Since there is a chance of UDP packet loss, and the goal for maximum packet loss is less than 1 in 10^6, a lightweight retry protocol will be used to ensure this aim.

The Linux cluster software stack and user workstation software suite comprise the control software that is being

* Mittuniversitetet, Sundsvall, Sweden

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

SODIUM LASER GUIDE STAR EMULATION

I. Price, RSAA, Australian National University, Australia
R. Conan, GMTO Corporation, Pasadena, CA 91101, USA

Abstract

In the era of extremely large telescopes (ELT) an adaptive optics (AO) system with artificial guide stars is an essential part of the optics between the source and the scientific instrument. In the case of the Giant Magellan Telescope these guide stars are formed by stimulating emission from Sodium atoms in the upper atmosphere with lasers launched from the side of the telescope. Moreover, they are resolved by the adaptive optics system so Shack-Hartmann wavefront sensors record elongated spots. Cost effective proof-of-concept systems for investigating control algorithms must be built for deployment in optics labs or on small telescopes. We present a hardware and software system that mimics the propagation of a single laser guide star (LGS) through the Earth's atmosphere and the optics of the Giant Magellan Telescope, using source motion and brightness modulation to simulate the source extension. A service oriented architecture allows adaptive optics scientists to construct images from different LGS asterisms and build non-real-time closed-loop control systems in high-level languages.

SODIUM LASER GUIDE STARS

A Sodium Laser Guide Star (LGS) is an artificial source created from Sodium atoms in the upper atmosphere and used for wavefront sensing as part of Adaptive Optics systems. Laser light at the 589 nm Sodium D2 line wavelength is propagated upward from a small telescope located alongside the main telescope. The photons from the laser interact with a naturally occurring population of Sodium atoms at an altitude of around 90-100km, exciting a shell electron. When the electron returns to its original energy level it emits a photon. Although the Sodium atoms are stimulated by a laser the process is dominated by spontaneous emission events. Ideally all of the energy from the outgoing laser beam would be delivered to atoms at a single altitude. In practice photons are absorbed all along the path through the Sodium layer, based on the remaining laser energy and the Sodium density. This results in a column of point sources that radiate some light back towards the main telescope, forming a laser guide star. Adaptive Optics systems use a wavelength filter to optically separate light from the laser guide star from the sources of scientific interest and natural guide star(s), and feed it to a laser guide star wavefront sensor (LGS WFS).

Shack-Hartmann wavefront sensors are based on lenslet arrays and imaging detectors such as CCDs. In essence, the pupil of the main telescope is divided into small patches, typically 100 to 200mm across. Each lenslet transforms the tip and tilt of the wavefront over the patch into a spot on the detector. Deviations in the spot positions gives an absolute measurement of the local tip and tilt of the wavefront. All lenslets in the WFS are observing the same LGS, but they each have a slightly different point of view, as illustrated in Figure 1. In the era of extremely large telescopes the size of the primary mirror is significant compared to a 15km long source 90km away. For a 25m diameter telescope, the spot elongation for the lenslet the furthest from the laser launch location will be 9.5". In contrast, a natural guide star spot has a typical size of 1".

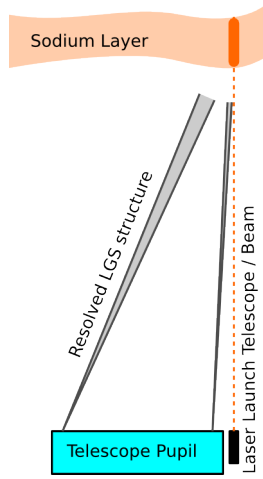


Figure 1: Column structure of the laser guide start source in the atmosphere and the variation of the angular size of the source with position in the telescope pupil.

Lenslets that sample the pupil far from the location of the laser launch telescope resolve the vertical structure of the LGS producing an elongated spot on the WFS detector, as shown in Figure 2. This distributes the signal from the LGS over more pixels than would be the case for an unresolved spot, increasing the error in spot position measurements and limiting the range of tip and tilt that can be measured without clipping of the input. Clearly the elongation of the spots has implications for the algorithms applied to WFS data for wavefront correction.

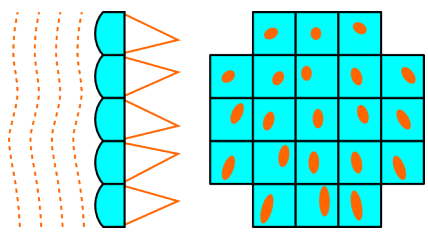


Figure 2: A distorted wavefront incident on the lenslet array of a Shack-Hartmann wavefront sensor (left) and the elongated spots in the image plane resulting from a laser launch telescope effectively located near the top edge of the pupil (right).

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

HOT CHECKOUT FOR 12 GeV AT JEFFERSON LAB*

R. Slominski, JLab, Newport News, VA 23606, USA
T. Larrieu, JLab, Newport News, VA 23606, USA

Abstract

A new hot checkout process was implemented at Jefferson Lab for the upgraded 12 GeV accelerator. The previous process proved insufficient in the fall of 2011 when a fire broke out in a septa magnet along the beam line due to a lack of communication about the status of systems. The improved process provides rigorous verification of system readiness thus protecting property while minimizing program delays. To achieve these goals, a database and web application were created to maintain an accurate list of machine components and coordinate and record verification checks by each responsible group. The process requires groups to publish checklists detailing each system check to encourage good work practice. Within groups, the process encourages two independent checks of each component: the first by a technician, and a second by the group leader. Finally, the application provides a dashboard display of checkout progress for each system and beam destination of the machine allowing for informed management decisions. Successful deployment of the new process has led to safe and efficient machine commissioning.

INTRODUCTION

Hot checkout (HCO) is the process by which all required accelerator hardware and software systems are tested for safe and effective accelerator operation prior to beam startup. Checks are performed while the machine is locked up and energized. The goal of HCO is to reduce program interruptions due to faulty or misconfigured hardware and software and to minimize the likelihood of a failure that damages equipment. Such an event occurred at Jefferson Lab (JLab) on October 14 2011 after a small number of preliminary 12 GeV upgrade changes had been put into place. Two septa magnets overheated and ignited a fire causing both program delays and property damage (Fig. 1). The low conductivity water valves for the magnets were closed and the thermal switches on the box supply that normally would trip on an over-temperature fault failed because the switch was disabled via a jumper. The root cause was determined to be inadequate communication of machine readiness and poor work control practices [1].

More extensive changes to the machine, and thus more complicated hot checkouts during the upcoming 12 GeV era were going to need strict oversight. At the time of the magnet fire there was no formally documented procedure for HCO. The previous approach relied on a single individual, the so-called Restoration Coordinator (RECO) to obtain verbal confirmation from the multiple group leaders of the readiness of their systems. This approach was appropriate for shorter downs where few changes were made to the machine as the RECO did not have the authority or resources to ensure the groups responsible for the machine were meticulously verifying component status. A committee was formed to address the shortcomings of the RECO and produce a more stringent process and associated software tool [2].



Figure 1: Septa magnet MYR7S03 fire damage.

PROCESS OVERVIEW

The new HCO process can be organized into three stages: setup, upkeep, and signoff. The initial setup stage involves documenting components and their properties and only has to occur once. The upkeep stage is on-going and ensures the data put in place during setup remains up-to-date. Upkeep tasks often coincide with accelerator maintenance periods and include preparation for each signoff period. The signoff stage entails checking component readiness and occurs immediately before resuming beam operations after an extended accelerator maintenance period. In practice the signoff stage is usually done twice a year: once before the fall run and again before the spring run.

Scope

The rigorous HCO is reserved for use following extended maintenance periods of a week or more The HCO process is not used during machine running to track downtime as the HCO process is not designed to be a repair log. There is a separate process and tool for tracking downtime at JLab during machine operation [3]. The HCO process is analogous to an aircraft pre-flight checklist process: once the plane is in the air the checklist is no longer useful.

*Authored by Jefferson Science Associates, LLC under U.S. DOE Contract No. DE-AC05-06OR23177

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

DRIFT CONTROL ENGINES STABILIZE TOP-UP OPERATION AT BESSY II*

T. Birke†, F. Falkenstern, R. Müller, A. Schälicke
Helmholtz-Zentrum Berlin für Materialien und Energie, Berlin, Germany

Abstract

Full stability of orbit and bunch-by-bunch-feedback controlled top-up operation becomes available to the experimental users only if the remaining slow drifts of essential operational parameters are properly compensated. At the light source BESSY II these are the transversal tunes as well as the path length and energy. These compensations are realized using feedback control loops together with supervising state machines. Key to the tune control is a multi source tune determination algorithm. For the path length correction, maintenance of beam energy and reduction of orbit deviations empirical findings are utilized.

INTRODUCTION

At lightsources, capabilities for user experiments rely on source point position, pointing stability as well as stable beam size and shape. These parameters transform to electron beam orbit and tune stability, typically controlled by slow (SOFB) and fast (FOFB) orbit correction feedback installations and by insertion device (ID) focusing feedforward compensations.

The intrinsic stability of the storage ring is greatly enhanced by the thermal equilibrium provided by top-up operation. In consequence correction amplitudes of beam guiding and shaping systems are drastically reduced, and beam quality can be better maintained by the described means.

Nevertheless, new constraints are imposed by the top-up injection permission, the technical solution implemented for FOFB and the stability requirement for efficiency of the Bunch by Bunch FeedBack (BBFB) [1] and for the amplitude of the resonantly excited bunch used for pseudo single bunch pulse picking. At BESSY II, additional control tools are needed to keep tunes, orbit and energy stable within required limits.

MOTIVATION

The ID-tune-feedforward system compensates the baseline tune shifts produced by gap and shift movements of any ID. This feedforward system interpolates empirically produced tables generated and iterated during measurements in dedicated ID-commissioning runs. Nevertheless imperfections and multiple IDs moving gaps and shifts at the same time would still cause significant tunes-shifts if not taken care of.

The purity of the camshaft bunch in the ion-clearing gap is controlled by a vertical knock out kicker. The resonant vertical beam excitation pulse during the gap around the camshaft bunch (see range marked blue in Figure 1) has to be very short to minimize perturbation of the camshaft bunch. So the kicker is active only for a few damping times at the injection shot. It has to be very efficient to prevent accumulation of charge in bunches to be cleared: If accumulated charge is cleared later when exact resonance condition is reestablished this might result in a shot injection efficiency below the allowed limits. Thus perfect matching of the excitation frequency and the vertical tune is needed to ensure excellent knockout efficiency.

Damping capability of the BBFB systems and the appropriate operational safety margin are well guaranteed as long as the tunes are constant.

Frequency and amplitude of the incoherent excitation for Pulse Picking by Resonant Excitation (PPRE) [2] have to match a specific frequency at the slope of the horizontal tune synchrotron sideband. Even tune shifts of about 500 Hz have severe impact on excitation effectiveness and stability of the PPRE bunch (see Figure 2).

Path length drifts were compensated by the slow orbit feedback (SOFB) by integrating the RF frequency into the horizontal response matrix and minimizing average horizontal corrector strength in the SVD calculation at each step. However, a straight forward implementation of this scheme was not possible with the fast orbit feedback system (FOFB).

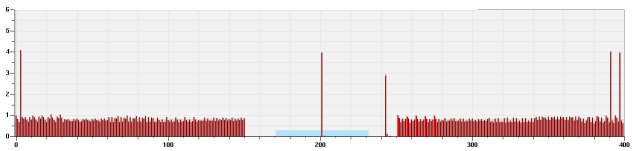


Figure 1: Fillpattern in standard multi- and single bunch hybrid mode, serving both average and timing mode experiments. A multibunch fill of 300 bunches including 3 bunches for fs-slicing experiments with an ion-clearing gap of 200 ns (100 buckets). A camshaft bunch at a fixed position in the center and a PPRE bunch close to the end of the gap - both for pseudo single bunch experiments. The blue region on both sides of the camshaft bunch is the cleared by resonant RF knockout during every injection shot.

TRANSVERSAL TUNE MEASUREMENT

Input Signals

Beammotion over a wide frequency range up to 1250 kHz is measured at a crossed stripline with a fast ADC/PXI system running LabVIEW. A processed FFT is provided to the control system in 3 partial ranges for each horizontal, ver-

* Work funded by BMBF and Land Berlin
† thomas.birke@helmholtz-berlin.de

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

EUROPEAN XFEL CAVITIES PIEZOELECTRIC TUNERS CONTROL RANGE OPTIMIZATION

W. Cichalewski, A.Napieralski, LUT-DMCS, Lodz, Poland
J. Branlard, C. Schmidt, DESY, Hamburg, Germany

Abstract

The piezo based control of the superconducting cavity tuning has been under the development over last years. Automated compensation of Lorentz force detuning of FLASH [1] and European XFEL [2] resonators allowed to maintain cavities in resonance operation even for high acceleration gradients (in range of 30 MV/m). It should be emphasized that cavity resonance control consists of two independent subsystems. First of all the slow motor tuner based system can be used for slow, wide range mechanical tuning (range of hundreds of kHz). Additionally the piezo tuning system allows for fine, dynamic compensation in a range of 1 kHz. In mentioned pulse mode experiments (like FLASH), the piezo regulation budget should be preserved for in-pulse detuning control. In order to maintain optimal cavity frequency adjustment capabilities slow motor tuners should automatically act on the static detuning component at the same time. This paper presents work concerning development, implementation and evaluation of automatic superconducting cavity frequency control towards piezo range optimization. FLASH and XFEL dedicated cavities tuning control experiences are also summarized.

INTRODUCTION

Both FLASH and European XFEL are free electron lasers facilities that build up accelerated beam energy using superconducting linacs. Superconducting cavities are or will be operated in pulse mode with 10 Hz repetition range and field gradients up to 30 MV/m. This work conditions causes extensive Lorentz force based reaction on the structure walls [3]. This cause mechanical deformation in range of few micrometers. For around 1 meter long, 1,3GHz resonant frequency, niobium resonator such a length change induces dynamic detuning modification in the range of couple hundreds of hertz. As cavities loaded quality factor is high (range from 3e6 for FLASH to 4.6e6 for XFEL) such misalignment results in significant accelerating field gradient drop. This have to be compensated by increase of supplying RF power - in order to maintain constant beam energy level. The other possibility to minimize this effect is external mechanical excitation provided by slow (step motors) and fast (piezo elements) cavity tuners.

MOTIVATION

Since fast tuners are foreseen for precise, in-pulse reaction to cavity dimension deviation theirs parameters have to be well recognized and adequately used.

Cavity Tuning Components Characterization

The European XFEL linac installation preparation includes also cryo-module testing phase [4]. Among the others also piezoelectric components are being characterized. One of test goal is to determine coefficient that describes relationship between piezo supply voltage and superconducting resonator frequency shift. The same study allows to verify both piezo elements (actuator and sensor) tuning range and voltage polarity (see Figure 1). Other (LLRF related) test

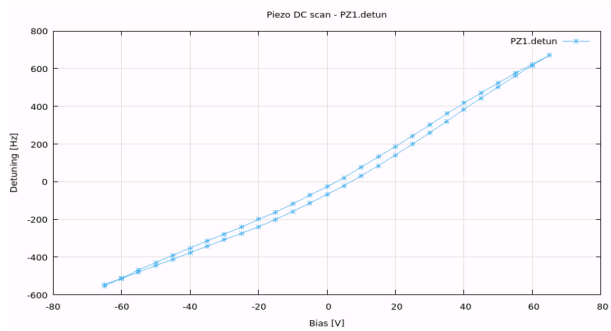


Figure 1: Piezo element tuning range characterization example (XFEL module XM51, cavity 3).

is dedicated to cavity slow motor tuners examination. This devices tuning range is not being verified (about hundreds of kHz). The cavity detuning in function of the motor position change is determined in order to evaluate transfer coefficient (like in previous test). Moreover motor backlash and movement hysteresis is determined. Exemplary measurement results can be found in Figure 2. All measurements results

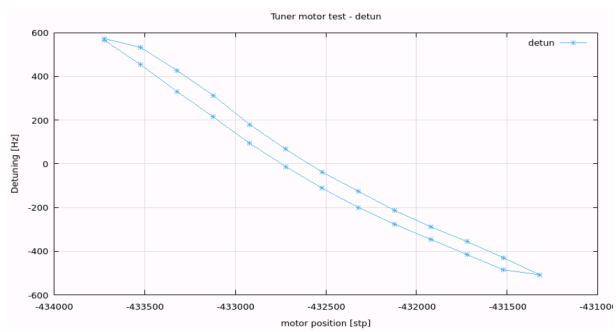


Figure 2: Slow tuner motor behavior characterization (XFEL module XM51, cavity 3).

are stored in dedicated (PostgreSQL [5] based) database. This data will be used during the XFEL linac commissioning for optimal LLRF system configuration. Additionally this information is also vital for described piezo relaxation algorithm configuration.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

CONTROL SYSTEM OF RF STATIONS FOR NICA BOOSTER

G. A. Fatkin, A. M. Batrakov, I. V. Ilyin, M.Yu. Vasilyev, BINP SB RAS, Novosibirsk, Russia
G. A. Fatkin, NSU, Novosibirsk, Russia

Abstract

NICA (Nuclotron based Ion Collider fAcility) is an accelerator complex, which is being built in JINR (Dubna, Russia). The system described in this paper is controlling the RF stations of Booster, the first element of the NICA complex. Control System consists of two devices: Intellectual Controller and Tester module. The first one is designed for precise measurement of magnetic field, generation of the acceleration frequency in accordance with measured field and control of RF power and pre-amplifiers. Intellectual Controller is a real-time feed-forward system with 40 μs loop time. It is based on ARM microcontroller and bare-metal control programs are used to reach maximum performance. Approaches that were used to achieve maximum performance are elaborated and presented in this paper. The second part of system - Tester is a simulator for tuning and checking the RF stations before start of operations or in absence of real accelerator. The achieved accuracy in chain 'magnetic field' – 'acceleration frequency' is better than $5 \cdot 10^{-5}$. Plans on feedback incorporation to stabilize ion beam behavior via frequency and phase tuning are discussed.

INTRODUCTION

The superconducting accelerator complex NICA (Nuclortron based Ion Collider fAcility) is constructed in JINR, Dubna. Project aims to provide collisions of heavy ion beams in the energy range from 1 to 4.5 GeV/u at the luminosity level of $1 \cdot 10^{27} cm^{-2} \cdot s^{-1}$. The Booster is a cycling accelerator of 197Au32+ ions and serves as one of the elements in the injection chain. It accepts particles with energy 6.2 MeV/u, accelerates them and extracts with energy 600MeV/u to the next acceleration stage - Nuclotron [1].

RF subsystem of the Booster consists of: two cavities, two RF stations, intellectual controller and tester module. RF stations include power amplification cascades and low-voltage electronics. They provide 10 kV acceleration voltage in required frequency range (0.5-5 MHz) on cavity gaps [2]. Primary function of the controller is to generate master frequency depending on the current value of magnetic field. Inaccuracy of frequency setting must not be worse than $\pm 2 \cdot 10^{-5}$.

RF CONTROL STRUCTURE

RF control structure is shown on Fig. 1. Signals from induction sensor and synchronization pulses are passed through the Tester module to Controller. Tester module allows to substitute real signals with imitated scenario or to through-pass these signals. Frequency, phase and cavity voltages are adjusted according to measured magnetic field and acceleration stage. Other technological parameters are measured and set, they are presented at table 1. Both RF Controller and Tester modules use telnet over Ethernet interface for interaction.

Table 1: Input and Output Signals of Controller

Signal	Channels	Rate	Resolution
Output signals			
Master frequency	2	25 kHz	24
V cavity	2	25 kHz	12
I anode	2	100 Hz	10
Input signals			
Field sensor	1	25 kHz	18
V cavity	2	25 kHz	12
Phase difference	1	25 kHz	12
Phase bias	1	25 kHz	12
Frequency bias	1	25 kHz	12
V preamplifier	2	100 Hz	12
I anode	4	1 kHz	12
V rectifier	6	1 Hz	12
V filament	2	1Hz	12
Synchronization	7	N/A	N/A

Controller Structure

Observing table 1 one may conclude that controller module must supply 4 fast (25 kHz) output channels and measure 6 signals with the same rate. Most calculations and data conversions will require floating-point arithmetic. It is also necessary to arrange data interchange between several independent measurement and control devices. E.g. to generate master frequency controller must receive data from measuring ADC, convert, correct and integrate the resulting value to attain field value then calculate frequency, apply necessary corrections, then convert it again and finally send this data to DDS. At the same time other signals must be changed and measured: cavity voltages, cathode currents, preamplifier voltages etc.

Probably betatron and synchrotron oscillations damping would be necessary for Booster. It is very hard to predict feedback parameters before Booster starts operation. Therefore we decided to implement feedback later using separate module which will provide frequency and phase bias analog signals.

Nowadays there are two main approaches for creation of such complex high rate systems: FPGA and high-end microcontrollers. Our reasons for choosing micro-controller were following: the need of floating point arithmetic, complex logic with different modes of operation that probably will require expansion and finally the need to provide TCP/IP interface for the interaction with Booster control system. At

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

TPS BOOSTER TUNE MEASURENET SYSTEM

P. C. Chiu, K. H. Hu, H. P. Hsueh, Y. S. Cheng, Demi Lee, K. T. Hsu
NSRRC, Hsinchu 30076, Taiwan

Abstract

The Taiwan Photon Source (TPS) is a state-of-the-art synchrotron radiation facility featuring ultra-high photon brightness [1]. Its Booster has 6 FODO cells which include 7 BD dipoles with 1.6 m long and 2 BH dipoles with 0.8 m long in each cell. After magnetization of stainless steel vacuum chamber of the booster were identified and then dismantled, annealed, and re-installed, the electron beam energy of the TPS booster ring has ramped to 3 GeV in a week. The booster tune correction during ramping is one of the main reasons why the booster commissioning progress so fast. This report will be summarized the booster tune monitor system.

INTRODUCTION

The TPS is a state-of-the-art synchrotron radiation facility featuring ultra-high photon brightness with low emittance. The TPS accelerator complex consists of a 150 MeV S-band linac, linac to booster transfer line (LTB), 0.15–3 GeV booster synchrotron, booster to storage ring transfer line (BTS), and 3 GeV storage ring. The booster has 6 FODO cells and its circumference is 496.8 meter; the storage ring's circumference is 518.4 meters with 24 DBA lattice and 6-fold symmetry. The booster and the storage ring share the same tunnel in a concentric fashion. During 4 years of construction period, civil constructions had been completed in early 2013. The accelerator installation then had taken another one year.

At September 2014, booster tune measurement commissioning had committed with beam commissioning. First-turn was easily obtained by beam steering and then multi-turn was also observed after optimization of Linac and transfer line. However, the beam could not be stored. After some hardware improvement such as power supply tuning, chamber and magnet re-alignment, kicker and septum improving and etc., the key setback was finally found on November 12: the pipes had a relative high permeability (ranging from 1.2 to 2.0) induced from the lack of proper annealing process. These un-annealed chambers were uninstalled and treated in vacuum oven up to 1050 ℃, and then re-installed. Booster then had stored beam in DC mode and ramped to 3 GeV successfully at December 16 2014 after several times of tune compensation. The tune measurement system provides precise tune measurement to correct tune variation during energy ramping and avoid across the resonant line which would cause a lot of beam loss. This report will summarize the infrastructure of booster tune measurement system measurement results and tune compensation during TPS booster commissioning.

TUNE MONITOR SETUP

Originally during booster commissioning beginning in Sep. 2014, the magnetic shakers where two multi-turn coils are mounted on vacuum chamber in horizontal and vertical plane are applied to excite beam. The kickers with 50Ω terminated load have calibration factor of 3 mG/A and are driven by a 50W amplifiers. Later, the stripline electrodes are adopted to replace magnet shakers on the booster synchrotron considering more power strength to excite beam. The TBT data provided by BPM electronics would be acquired to extract tune by FFT. Agilent arbitrary signal generator would provide band-limited, strength-adjustable excite signal. The functional block diagram of this new tune monitor system is shown in Fig.1. Figure 2 shows the stripline kicker installed in June 2015 to provide more efficient power to excite beam.

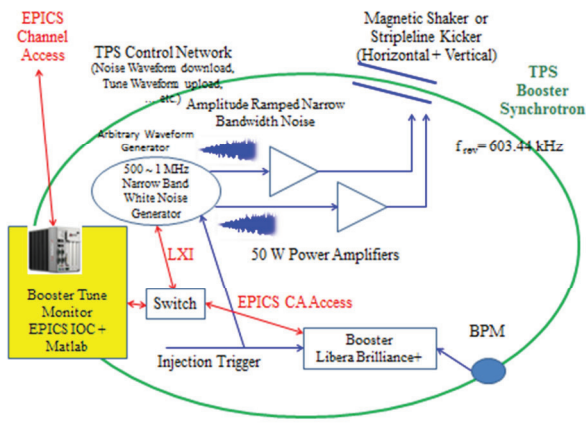


Figure 1: Functional block diagram of the tune monitor for TPS booster.

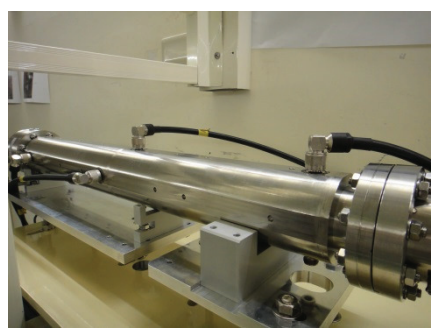


Figure 2: Stripline kickers installed in June 2015.

Figure 3 shows the booster tune measurement GUI. The spectrogram of the horizontal and vertical of BPM turn-by-turn data could identify tune variation clearly. Peak identification from the spectrogram could extract the

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

INTEGRATING THE MEASURING SYSTEM OF VIBRATION AND BEAM POSITION MONITOR TO STUDY THE BEAM STABILITY

C. H. Huang, C.Y. Liao, Y.S. Cheng, P. C. Chiu, K. H. Hu, K. T. Hsu
NSRRC, Hsinchu 30076, Taiwan

Abstract

For a low emittance, third-generation light source such as Taiwan Photon Source (TPS), beam orbit motion needs to be controlled within submicron for obtaining a high quality light source. Magnets vibration especially quadruples will be one of the main sources to destroy the beam stability. In order to study the relationship between vibration and beam motion, a synchronous data acquisition system which integrates the measurement of vibration and beam position monitoring system is highly desirable. For a larger vibration such as earthquakes are also deleterious to beam stability or even make the beam trip due to the quench of superconducting RF cavity. A data acquisition system integrated with an earthquake P-wave detector is also quite necessary to show and archive the data on the control system. First, the data acquisition systems of vibration and earthquake measurement system are summarized in this report. The preliminary study of the relationship between the beam motion and magnets vibration will also be presented here.

INTRODUTION

For a low emittance, third-generation light source such as TPS, it imposes a stringent requirement for the orbit stability. To obtain a high quality light source, the beam orbit motion needs to be controlled at least within 0.5 μm. There are many sources [1] which may destroy the beam stability such as vibration, ill function of power supply, septum/kicker field leakage, etc. The ground vibration and the technical noise such as vacuum pumps and cooling water flow will cause the magnet and vacuum chamber vibration. It leads to the distortion of the close orbit [2]. The dominant effect for the beam motion caused by magnet vibration is produced by the quadrupole and it introduces a kick angle to the beam as the quadrupole has an offset to the beam due to the vibration.

To achieve submicron stability, various efforts and studies should be continuously performed such as power supply performance, RF system performance, cooling water, mechanical vibration, orbit feedback system and bunch-by-bunch feedbacks. Besides, it is also important to identify various sources and to minimize their effects to the beam instability. In this paper, the data acquisition system for measuring the vibration and beam position is introduced first. The setup of beam position monitoring and earthquake monitoring system are described in the third and fourth parts. In the fifth and sixth sections, the preliminary results to identify the vibration caused by turbo molecular pump (TMP) system and to study the beam stability are shown. Finally is the short summary.

VIBRATION DETECTOR AND DATA ACQUISITION SYSTEM

Low noise three-component seismometers, LE-3Dlite Mark II, with frequency range 1- 100 Hz is used in this study. The data acquisition unit (Data Translation DT8837), which is complied with LXI class C standard, provides Ethernet accesses via SCPI command to acquire data. Multiple DT8837s are synchronized by wired trigger bus (WTB) interface. To extend length limit of WTB cable, an in-house made small interface adapter from RJ-45 to Micro D is installed at the WTB connector of the DT8837 side. It allows unshielded twisted pair (UTP) cables to replace WTB cables and to send the trigger, sync and clock single from the timing system adapter to DT8837 more than 100 m. Coherent data acquisition can also be achieved by the aid of global timing system for long rang or by LXI WTB cable for short range application. A Matlab script can be running in the server for the long-term measurement or in the laptop for the short-term measurement, shown in Fig. 1. The synchronous data acquisition between vibration measurement and beam position monitoring system can also be through the timing system.

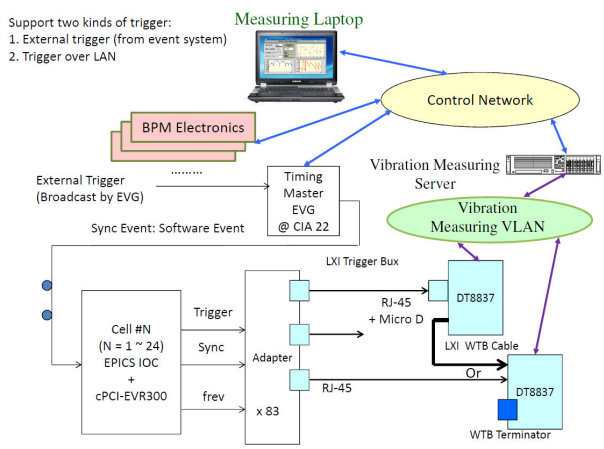


Figure 1: Configuration of vibration measuring system integrated with beam position monitoring system.

BEAM POSITION MONITORING SYSTEM

The TPS storage ring is divided into 24 cells and 7 beam position monitors (BPMs) are installed in each cell.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

CONTROL OF FAST-PULSED POWER CONVERTERS AT CERN USING A FUNCTION GENERATOR CONTROLLER

R. Murillo-Garcia, Q. King, M. Magrans de Abril, CERN, Geneva, Switzerland

Abstract

The electrical power converter group at CERN is responsible for the design of fast-pulsed power converters. These generate a flat-top pulse of the order of a few milliseconds. Control of these power converters is orchestrated by an embedded computer, known as the Function Generator/Controller (FGC). The FGC is the main component in the so-called RegFGC3 chassis, which also houses a variety of purpose-built cards. Ensuring the generation of the pulse at a precise moment, typically when the beam passes, is paramount to the correct behaviour of the accelerator. To that end, the timing distribution and posterior handling by the FGC must be well defined. Also important is the ability to provide operational feedback, and to configure the FGC, the converter, and the pulse characteristics. This paper presents an overview of the system architecture as well as the results obtained during the commissioning of this control solution in CERN's new Linac4.

INTRODUCTION

CERN has nine accelerators, numerous experiments and various test areas. One commonality of such infrastructure is the need of a power system to energize the magnets (dipoles, quadruples, septums, etc.) and the various radio frequency cavities. In total, CERN has around 5,600 power converters, accounting for a typical consumption during exploitation of 1.2 TWh per year [1].

From an operational point of view, power converters fall into three categories:

- Cycling: the field, current or voltage reference value can synchronously be modified on a cycle-by-cycle basis. At CERN each cycle is a multiple of a basic period, which has a fixed duration of 1.2 seconds.
- DC: the reference is asynchronously modified as required during operation to reach a desired field, current or voltage.
- Aperiodic: operators trigger aperiodic cycles to ramp the current to maintain the required beam stability. This is the case of the LHC.

Fast-pulsed converters are used with cycling circuits to generate a pulse with a flat-top duration of only a few milliseconds and synchronous to an event such as beam injection or extraction.

This paper describes the use of the third generation of the CERN-designed Function Generator/ Controller embedded platform (FGC3) to control the four types of fast-pulsed power converters currently being commissioned in CERN's new Linac4 (see Table 1).

The Linac4 beam is 400 µs long, whilst the flat-top pulse is stable solely during a few milliseconds. As these two events must be synchronized to ensure the correct functioning of the accelerator, the FGC3 must guarantee time accuracy and inter-FGC precision better than 10 µs.

Table 1: Fast-pulsed Power Converters in Linac4

Name	Peak Power	Pulse	Number*
Mididiscap	30 kW	5 ms	50
Maxidiscap	4 kW	2 ms	48
Modulator	5.5 MW	1.8 ms	14
HMinus	150 kW	1.2 ms	3

FUNCTION GENERATION/CONTROLLER

Most power converters currently being installed at CERN are controlled with FGC3s, shown in Fig. 1. This embedded device includes a mainboard, an Ethernet-based communication card† and an analog card with four high precision ADC channels (ADS1274) and two 16-bit DACs (MAX5541) [2]. The mainboard includes a Renasas RX610 microcontroller (running a small footprint real-time operating system called NanOS), a TI TMS320C6727 floating point DSP with a 10 kHz interrupt-driven task and a Xilinx FPGA for glue logic and peripheral handling.

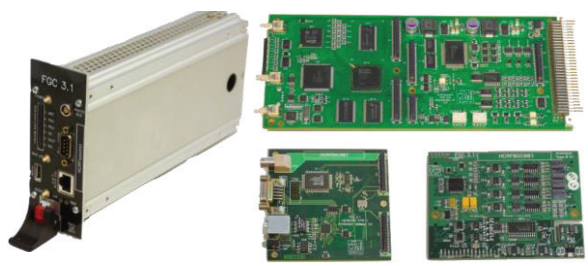


Figure 1: Function Generator Controller: mainboard, Ethernet and analog cards.

The FGC3 is integrated within a CERN-designed chassis called RegFGC3 [3], shown in Fig. 2. A backplane links the FGC3 with a variety of cards dedicated to state control, analog interlock, digital interlock, measurement conditioning, etc. Communication buses over this backplane include:

- SPIVS: serial bus used by the FGC3 to send the reference value in digital form.
- SCIVS: serial bus used by the FGC3 to exchange parameters with the cards.

* Number of converters on completion of Linac4.

† The FGC3 can also be fitted with a WorldFIP-based communication card.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

WHITE-RABBIT BASED REVOLUTION FREQUENCY PROGRAM FOR THE LONGITUDINAL BEAM CONTROL OF THE CERN PS

D. Perrelet, Y. Brischetto, H. Damerau, D. Oberson*, M. Sundal#, A. Villanueva, CERN, Geneva, Switzerland

Abstract

The measured bending field of the CERN Proton Synchrotron (PS) is received in real-time by the longitudinal beam control system and converted into the revolution frequency used as set-point for beam phase and radial loops. With the renovation of the bending field measurement system the transmission technique is changed from a differential sequence of pulses, the so-called B-train, to a stream of Ethernet frames based on the White Rabbit protocol. The packets contain field, its derivative and auxiliary information. A new frequency program for the conversion of the bending field into the revolution frequency, depending also on parameters like radius of the accelerator and the particle type, has been developed. Instead of storing large conversion tables from field to frequency for fixed parameters, the frequencies are directly calculated in programmable logic (FPGA). To reduce development time and keep flexibility, the conversion is processed in real-time in the FPGA using Xilinx floating-point primitives mapped by a higher level tool, Simulink System Generator. Commissioning with beam of the new frequency program in the PS is progressing.

INTRODUCTION

The CERN PS accelerates protons and ions [1] for fixed target experiments and as part of the LHC injector chain. For protons, its kinetic range covers 1.4 GeV at injection to 26 GeV total energy. The longitudinal beam control system drives in total 25 cavities in a wide frequency range (2.8-10, 20, 40, 80 and 200 MHz). It requires the so-called open-loop revolution frequency to be derived from the bending field of the main dipole magnets which serves as a reference for the feedback loops closed around the beam. In addition to protons from the PS Booster (PSB) the PS accelerates a large variety of heavy and light ions such as $^{208}Pb^{54+}$ [2,3,4] or $^{40}Ar^{11+}$ injected from LEIR.

Motivated by the renovation of the present measurement system of the magnetic field, its distribution in the form of asynchronous pulses (so-called B-train) [5, 6] has been upgraded [7] to a stream of Ethernet frames based on the White Rabbit (WR) protocol. The reception of the magnetic field information for the longitudinal beam-control and its conversion into the open-loop revolution frequency has been redeveloped and commissioned with beam.

*HEIA-FR (Haute école d'ingénierie et d'architecture Fribourg)
#IST (Instituto Superior Técnico Lisboa)

FIELD MEASUREMENT

The existing system uses a peaking-strip (PKS) marker at 49.8 Gauss as an absolute reference at the start of each cycle. This implies to ramp down to a low field between cycles to reset the integration process, consuming time and power. Moreover, the B_{up}/B_{down} pulses distribution scheme triggered by 0.1 G changes in B at a maximum rate of 500 kHz is at the limit for ions in terms of resolution and also in case of flat zones.

In the new system, the measurement setup is referenced at each cycle by a Ferromagnetic Resonance (FMR) tunable marker and integrates a flux coil voltage, meaning the derivative of B. This kind of marker has a working point closer to the injection field, thus reducing the ramp down constraint besides improving the injection accuracy. The regulation done in B, has to cover the range from 600 G to 13000 G with a maximum changing rate of ~ 30 kG/s.

Following the acquisition process, the magnetic field and its derivative are transported via Ethernet frames using the White Rabbit protocol over optical fibres [8]. At the reception side, B is extracted and then converted into the open-loop revolution frequency.

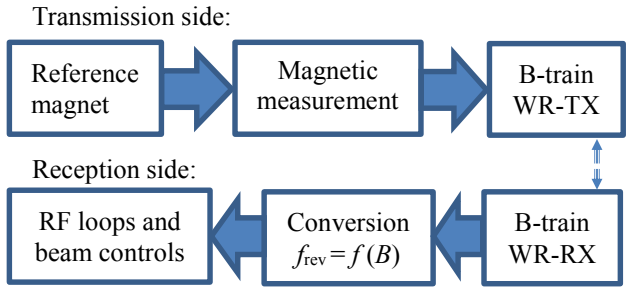


Figure 1: Overview of the whole system and its connection between reference magnet and longitudinal beam control.

OPEN-LOOP REVOLUTION FREQUENCY

In a synchrotron, the orbit of the particles can be controlled by changing the bending field in the dipole magnets or the RF frequency. A theoretical relationship between magnetic field and revolution frequency can be derived from machine parameters [9,10]. According to

$$f_{rev}(B,b) = \frac{c}{2\pi R_{nom}} \frac{1}{\sqrt{1+\left(\frac{E_0}{b \cdot B \cdot c \cdot \rho_{nom}}\right)^2}} \quad (1)$$

with the velocity of light in vacuum c, the mean radius R_{nom}

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

INTEGRATION OF THE TRACK BEAM DYNAMICS MODEL TO DECREASE LINAC TUNING TIMES

Christopher E. Peters#, Clayton Dickerson, Francisco Garcia, Maria A. Power
Argonne National Laboratory, Argonne, IL, USA

Abstract

The Accelerator R&D Group within the Argonne National Laboratory (ANL) Physics Division maintains a beam dynamics model named TRACK. This simulation code has the potential to assist operators in visualizing key performance parameters of the Argonne Tandem Linear Accelerating System (ATLAS). By having real-time access to visual and animated models of the particle beam transverse and longitudinal phase spaces, operators can more quickly iterate to a final machine tune. However, this effort requires a seamless integration into the control system, both to extract initial run-time information from the accelerator, and to present the simulation results back to the users. This paper presents efforts to pre-process, batch execute, and visualize TRACK particle beam physics simulations in real-time via the ATLAS Control System.

INTRODUCTION TO ATLAS AND TRACK

The ATLAS accelerator is located at the United States Department of Energy's Argonne National Laboratory in the suburbs of Chicago, Illinois. It is a National User Facility capable of delivering ions from hydrogen to uranium for low energy nuclear physics research in order to perform analysis of the properties of the nucleus. As a result of the wide variation of beams delivered [1], re_ tuning of the entire machine is necessary on a near weekly basis. After a series of upgrades, ATLAS will consist of two possible source lines, a common injection and beam transport line, and 8 different target areas. This wide range of possible machine configurations combined with the thousands of individual devices which support them present a very real challenge to operators to arrive at a final tune quickly.

The TRACK beam simulation program developed at ANL-PHY is available and has been validated during the design of upgrade projects at ATLAS [2]. This code can perform individual particle simulations of a beam traveling through various types of optical and RF fields in order to predict final emittance and beam shape. The resultant plots show an animation in time of beam emittance ellipses, and an overall plot of beam size vs. distance. It should be possible to link the TRACK simulation code to the real-time running parameters of ATLAS in order to give the operators a detailed overview of the tune process. This effort requires that the program be integrated into the control system in a seamless way so that non-PhD personnel can operate and understand the behavior of the machine. The TRACK program's graphical user interface (GUI) only runs in Windows (Figure 1) while the control system is primarily developed using Vista Control Systems, Inc.'s Vsystem [3].

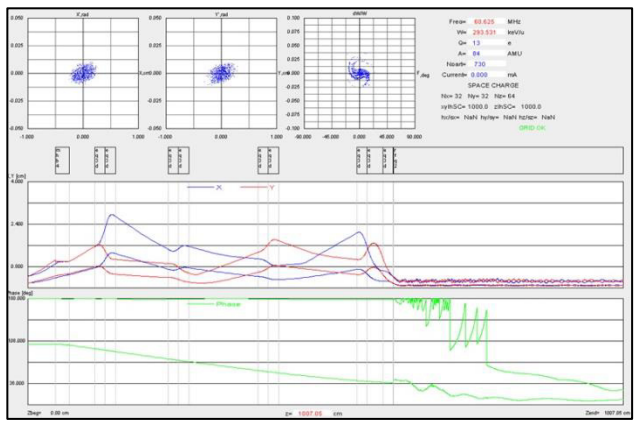


Figure 1: Existing TRACK GUI and Simulation.

USER INTERFACE & CONFIGURATIONS

In order to fully realize the potential of fast, real-time iterations on tuning time, the interface to TRACK must operate very similarly to the existing control system. Therefore, brand new interface windows were designed using native control system libraries (Figure 2). One important difference between the TRACK application GUI and the control system GUI is the presence of 'configuration' buttons. These are pre-selected input conditions, source/target line selections, and graph views which can be executed more quickly than re-running a full simulation for the entire machine.

Figures 3 and 4 represent a full TRACK simulation from current device settings by the operator. Note the black boxes which indicate corresponding quadrupole devices. By using the same interface as normal, operators can move sliders and immediately see the resultant simulation update.

USE CASES OF TRACK AT ATLAS

There are two main ways in which a pre-processed simulation of the beam could be useful to operators. The first is by importing the real-time settings of each device in the current beamline to give operators an overview of the performance of the machine. The second is by utilizing a database of previous runs to predict which machine configurations have been successful in the past. This second method will be useful before the start of a new experiment to predict future good tunes.

This work was supported by the U.S. Department of Energy, Office of Nuclear Physics, under Contract No. DE-AC02-06CH11357. This research used resources of ANL's ATLAS facility, which is a DOE Office of Science User Facility.

#ChrisPeters@anl.gov

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

REAL-TIME BEAM LOADING COMPENSATION FOR SINGLE SRF CAVITY LLRF REGULATION

I. Rutkowski, M. Grzegrzółka WUT, Warsaw, Poland
Ł. Butkowski, C. Schmidt, DESY, Hamburg, Germany
M. Kuntzsch, HZDR, Dresden, Germany

Abstract

Stable and reproducible generation of a photon beam at Free Electron Lasers (FELs) necessitates a low energy spread of the electron beam. A low level radio frequency (LLRF) control system stabilizes the RF field inside accelerating modules. An electron beam passing through the cavity induces a voltage proportional to the charge and the cavity shunt impedance. The feedback loop tries to compensate for the perturbation after the accelerating gradient drops. The delay and high gain result in an overshoot or oscillations during transients. A feed forward signal can be applied to act on the plant simultaneously with the RF feedback. It can be generated off-line based on system characteristics and beam parameters or on-line using information obtained from the beam diagnostic systems. In the latter scheme fluctuations of the beam current are accounted for in real-time using an open-loop (feedforward) controller. The bunch charge detection scheme and its implementation is described. This paper describes results of the tests performed on the ELBE (Electron Linac for beams with high Brilliance and low Emittance) radiation source at the HZDR (Helmholtz-Zentrum Dresden-Rossendorf) facility using a MTCA.4-based LLRF control system.

INTRODUCTION

Stable and reproducible generation of a photon beam at Free Electron Lasers (FELs) necessitates a low energy spread of the electron beam. A low level radio frequency (LLRF) control system stabilizes the RF field inside accelerating modules. An electron beam passing through the cavity induces a voltage proportional to the charge and the cavity shunt impedance. For an SRF cavity with a high loaded quality factor (Q_L) compensation of the beam loading may require more RF power to be provided than is required to sustain the field gradient in the cavity, making it a major source of distortion of the RF field. Waveforms from a MTCA.4-based digital LLRF system using only a PI feedback controller are shown in Fig. 1 (gain in arbitrary units).

The doubling time of the integral controller is significantly longer than the pulse length (4 ms). Increasing the gain decreases the static error during the beam pulse, but results in transient oscillations (see Fig. 2). A feed forward signal can be applied to act on the plant simultaneously. It can be generated off-line based on system characteristics and expected beam parameters or on-line using information obtained from the beam diagnostic systems. It has been demonstrated [1] that real-time beam loading compensation

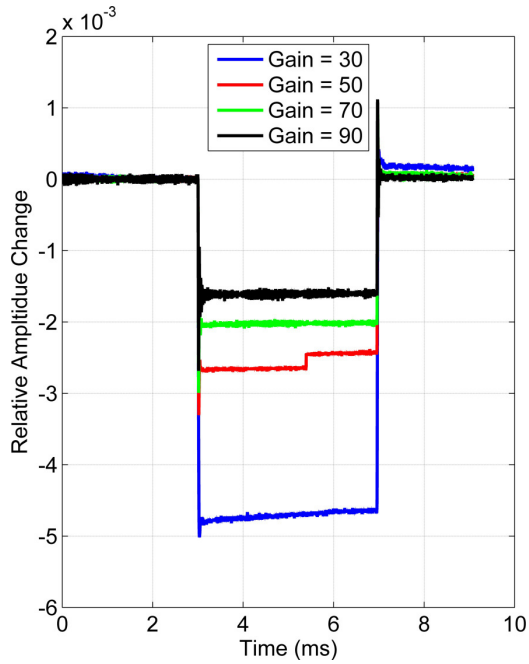


Figure 1: Response of a LLRF system using only a PI feedback controller to a beam pulse.

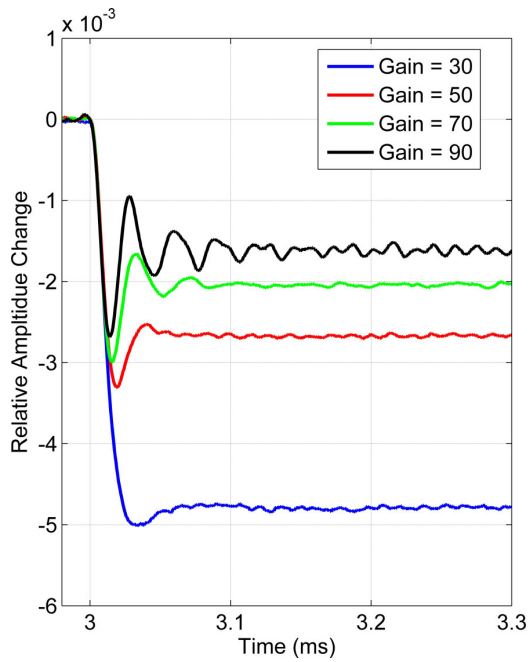


Figure 2: Response of a LLRF system using only a PI feedback controller to a beam pulse (zoom in during beam rise).

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

ARCHITECTURE OF TRANSVERSE MULTI-BUNCH FEEDBACK PROCESSOR AT DIAMOND

M.G. Abbott, G. Rehm, I.S. Uzun, Diamond Light Source, Oxfordshire, UK

Abstract

We describe the detailed internal architecture of the Transverse Multi-Bunch Feedback processor used at Diamond for control of multi-bunch instabilities and measurement of betatron tunes. Bunch by bunch selectable control over feedback filters, gain and excitation allows fine control over feedback, allowing for example the single bunch in a hybrid or camshaft fill pattern to be controlled independently from the bunch train. It is also possible to excite all bunches at a single frequency while simultaneously sweeping the excitation for tune measurement of a few selected bunches. The single frequency excitation has been used for continuous measurement of the beta-function. A simple programmable event sequencer provides support for up to 7 steps of programmable sweeps and changes to feedback and excitation, allowing a variety of complex and precisely timed beam characterisation experiments including grow-damp measurements in unstable conditions and programmed bunch cleaning. Finally input and output compensation filters allow for correction of front end and amplifier phasing at higher frequencies.

INTRODUCTION

At Diamond Light Source (DLS) we have been using a Transverse Multi-Bunch Feedback (TMBF) system based on the Libera hardware platform [1] for nearly a decade [2–5] to control multibunch instabilities. With a new bunch position every 2 ns it is appropriate for the feedback processing to be implemented in an FPGA. The original firmware was based on work done at the ESRF [6], but subsequently the firmware and control system have been rewritten to provide further extensions to the system functionality [7, 8]. This work has extended the lifetime of the original platform and increased its capability, and the DLS TMBF system has now been adopted at ALBA [9] and is being evaluated at the Australian Synchrotron.

In this paper we expand on and update the original IBIC 2013 [7] paper with many system refinements.

OVERVIEW

Figure 1 shows the TMBF processors in context. A 4-button electrode assembly together with an RF hybrid circuit picks up the horizontal and vertical position of each bunch. An RF front end converts this raw 3rd harmonic signal into a 250 MHz bandwidth bunch position signal suitable for processing by TMBF. This signal is digitised at 500 MHz (using four 125 MHz ADCs phased at 2 ns intervals), processed by the FPGA, and then reconverted to drive the striplines.

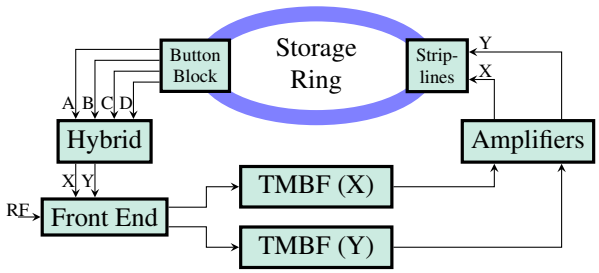


Figure 1: TMBF system in context. The Transverse Multi-Bunch Feedback system measures the position of each bunch, detects the betatron oscillations of each bunch, and generates a drive signal to suppress the oscillations.

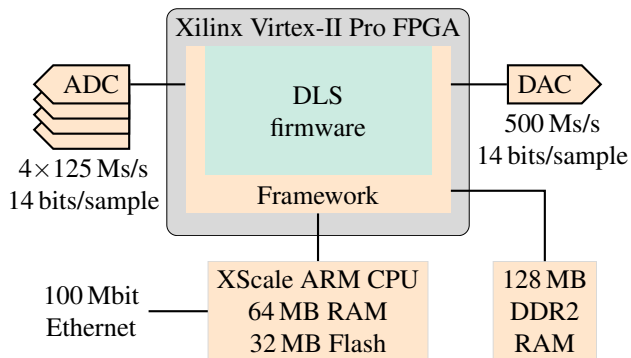


Figure 2: Libera System Platform. The DLS TMBF system is implemented on the Instrumentation Technologies Libera platform with the control system running EPICS on an ARM based embedded Single Board Controller (SBC).

The position of each bunch oscillates at the machine betatron tune, and these oscillations can normally be suppressed by feeding them back with a phase shift of 180°. The TMBF measures the horizontal or vertical position of each bunch, runs a 10-tap FIR filter on each bunch position to compute the appropriate phase adjustment, and outputs a negative feedback signal.

The TMBF processing system consists of analogue to/from digital converters connected to an FPGA for the high speed processing, together with memory for data capture and an embedded Single Board Computer as shown in Fig. 2.

As well as the core feedback function, a number of diagnostic functions are provided, both to enable easy monitoring of the system status, and to support complex experiments on the beam. System status monitoring includes detailed overflow detection and quick measurement of beam movement. Complex experiments are supported by internal oscillators, detectors, and a programmable sequencer, described in more detail below.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

PANDA MOTION PROJECT - A COLLABORATION BETWEEN SOLEIL AND DIAMOND TO UPGRADE THEIR 'POSITION AND ACQUISITION' PROCESSING PLATFORM

I.S. Uzun, T. Cobb, A. Cousins, M. Heron, Diamond Light Source, Oxfordshire, UK
YM. Abiven, J. Bisou, P. Monteiro, G. Renaud, SOLEIL, Gif-sur-Yvette, France

Abstract

Synchrotron SOLEIL and Diamond Light Source are two third generation light sources located respectively in France and the UK. In the past 5 years, both facilities separately developed their own platform permitting encoder processing to synchronize motion systems and acquisition during experiments, SPIETBOX by SOLEIL and ZEBRA by Diamond. New operational requirements for simultaneous and multi-technique scanning, and support of multiple encoder standards have been identified by both institutes. In order to address these a collaborative project has been initiated between SOLEIL and Diamond to realize a new 'Position and Acquisition' processing platform, called PandA. The PandA project addresses current systems' limitations in terms of obsolescence and need for more processing power. Its design is going to be a 1U standalone system powered by a Xilinx Zynq SoC to implement a configurable set of logic functionalities. It will provide a flexible and open solution to interface different third party hardware (detectors and motion Controllers). This paper details the organization of this collaboration, sharing technical leadership between both institutes and the status of the project.

COLLABORATION MANAGEMENT

SPIETBOX [1] and ZEBRA [2] were designed in-house by SOLEIL and Diamond respectively, and both have been integral part of beamline operations, and proven useful for synchronising various equipment and for position capture. One application is to synchronize detectors with optics that select the photon energy of the beam to acquire only useful data, another one is facilitating the use of continuous scanning for faster acquisition in 2014. Both institutes share the same vision for a new motion control and data acquisition architecture. So SOLEIL and Diamond have gone into a collaboration to re-design and improve these platforms with following guidelines:

- Both institutes want to standardize hardware, in order to reduce number of systems which need to be developed and maintained by the support staff.
- Develop a ready to use and complete solution, in order to install quickly and easily all the hardware and software components.

The main goal of the collaboration was to share resources and experiences between both synchrotrons for the development. The first step was to understand existing skills and human resource availability at both institutes and agree on sharing the responsibilities of each task at design level as shown in Table 1. We decided to manage the work through regular video conference meetings every two weeks, face-to face meeting every 3 months, as well as sharing minutes and documentation over a private Github repository. Once the organization was set-up, a set of hardware, firmware and software specifications were defined as described in the following sections. It is agreed that the hardware developed during the course of the collaboration would be developed under the CERN Open Hardware (OHW) Licence [3] and released on the public OHW repository following completion of the project. Both parties would also have the option to engage third parties to commercialise the resulting project.

Table 1: Task Sharing Between Project Partners

Task leadership	DIAMOND	SOLEIL
Project management	X	
HW, FW specifications	X	X
Schematic Design		X
PCB Layout		X
Mechanics		X
FPGA – Zynq Processor Design	X	
FPGA – Zynq Logic Design	X	
Linux Kernel Development	X	
Linux Application Development	X	
TANGO interface		X
EPICS interface	X	

HARDWARE ARCHITECTURE

New applications [4], operational feedback from beamline scientists and from the Flyscan project [5] identified the following requirements to address limitations of existing SPIETBOX and ZEBRA solutions:

- Multi Channel TTL and LVDS I/Os for synchronous triggering and clocking with higher timing resolution.
- 4-Channels of encoder interface I/Os, supporting a wider range of protocols including Incremental, Absolute, EnDat and BiSS.
- 2-Channels of SFP Gigabit Transceiver interface for IpBus, Timing System, Diamond Communication Controller [6] or custom high-speed serial connectivity.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

UPGRADED CONTROL SYSTEM FOR LHC BEAM-BASED COLLIMATOR ALIGNMENT

G. Valentino*, G. Baud, M. Gasior, S. Jackson, L. Jensen, J. Olexa, S. Redaelli, J. Wenninger
CERN, Geneva, Switzerland

Abstract

In the Large Hadron Collider (LHC), over 100 movable collimators are connected to a three-tier control system which moves them to the required settings throughout the operational cycle from injection to collision energy. A dedicated control system was developed to align the collimators to the beam during machine commissioning periods and hence determine operational settings for the active run. During Long Shutdown 1, the control system was upgraded to allow beam-based alignments to be performed using embedded beam position monitors in 18 newly installed collimators as well as beam loss monitors. This paper presents the new collimation controls architecture for LHC Run II along with several modifications in the Java-based application layer.

INTRODUCTION

The Large Hadron Collider (LHC) accelerates two counter-rotating beams to an energy of 7 TeV before colliding them in four points where experiment detectors are located [1]. Uncontrolled beam losses can result in quenches of the super-conducting magnets, which can damage the machine. For this reason, a beam collimation system is installed in the LHC, comprising over 100 collimators [2]. It is also designed to protect the LHC in the event of fast failures, which might damage accelerator components.

Each collimator installed for ring cleaning consists of two jaws, made of graphite, tungsten or copper, which have to be positioned around the beam at all times with an accuracy of less than 50 μm. The operational settings of the collimators are defined in order to establish a four-stage hierarchy. They are determined via a beam-based alignment procedure [3] which uses feedback from a Beam Loss Monitoring (BLM) [4] detector positioned downstream from the collimator. From the aligned jaw positions, the measured beam centers and beam sizes can be calculated.

The primary collimators (TCP) are positioned closest to the beam, followed by the secondary collimators (TCSG), tertiary collimators (TCT) and absorbers (TCLA). Most of the collimators are installed in Insertion Region (IR) 3 and 7 for off-momentum and betatron cleaning respectively. The TCTs are installed upstream of the experiment detectors (IR1, 2, 5 and 8), and a TCSG is installed in the dump region (IR6).

The software architecture for the LHC collimation system respects the standards of the LHC Software Architecture (LSA) [5], which consists of a 3-tier structure. For the collimators, the bottom layer is composed of a double PXI system that is used to control and read out stepping motors, sensors and measurement devices. The collimator jaw positions can be positioned with an accuracy of 5 μm, i.e. less than 2% of the 1σ beam size at the primary collimators at 7 TeV. The maximum jaw movement rate is 2 mm/s.

Application servers that host databases and operational files make up the middle layer, on top of which Graphical User Interface (GUI) console applications run. The collimator positions can be set from a remote location in the CERN Control Centre (CCC) and synchronized with the operation of the LHC cycle. The software applications interact with the hardware via the Common Middleware (CMW) [6] and Front-End Software Architecture (FESA) [7] infrastructures.

During Long Shutdown 1 (LS1), all 16 TCTs and the 2 TCSGs in IR6 were replaced by a new design, in which Beam Position Monitor (BPM) pick-ups are embedded in the upstream and downstream corner of each collimator jaw [8, 9]. The LHC installation followed a successful validation of a prototype with beam in the Super Proton Synchrotron (SPS) [10]. This required an upgrade of the beam-based alignment control system to acquire and use both the BLM and the BPM data for the alignment.

ALIGNMENT PROCEDURES

BLM-based Alignment

A four-stage procedure is used to align the collimators with feedback from the BLMs. Both jaws of a reference collimator are moved towards the beam in steps of 5-20 μm (1) until they reach the beam halo on either side, which is established when an appropriate loss pattern is observed in a downstream BLM. Then, the collimator for which the beam center and beam size need to be measured is aligned (2). The center is given as the average of the two aligned jaw positions. The reference collimator is then re-aligned (3), and the beam size at the previous collimator is given as the ratio of its gap in mm when aligned to the average of the cut in units of beam σ. The final step is to open the collimator to the new operational positions (4). This procedure is performed at very small gaps of $< 4\ \sigma$.

One of the main drawbacks of this procedure is that it causes losses of a fraction of the circulating beam and therefore can only be carried out with safe beam intensities. Collimator alignment can therefore not be performed dur-

* gianluca.valentino@cern.ch

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

HIGH LEVEL CONTROLS FOR THE EUROPEAN XFEL

Lars Fröhlich, Bolko Beutner, Winfried Decking, Olaf Hensler, Raimund Kammering,
Torsten Limberg, Sascha Meykopff, Josef Wilgen, DESY, Hamburg, Germany

Abstract

The European X-Ray Free-Electron Laser (XFEL) will generate extremely short and intense X-ray flashes from the electron beam of a 2.1 km long superconducting linear accelerator. Due to the complexity of the facility and the sheer number of subsystems and components, special emphasis needs to be placed on the automatization of procedures, on the abstraction of machine parameters, and on the development of user-friendly high-level software for the operation of the accelerator. The paper gives an overview of the ongoing work and highlights several new tools and concepts.

INTRODUCTION

The European X-ray Free-Electron Laser (XFEL) is a research facility currently under construction in close collaboration between the European XFEL Facility GmbH[1] and DESY[2] in Hamburg, Germany [1–3]. It consists of a superconducting linear accelerator delivering an electron beam with particle energies up to 17.5 GeV and a total beam power up to ~600 kW, several long undulator sections enabling the generation of extremely brilliant X-ray pulses at wavelengths down to 0.05 nm, beamlines and setups for photon science experiments, and the associated infrastructure. An overview of the control system architecture is given in [4].

Compared to other accelerator facilities operated by DESY, the XFEL presents itself as a system of unprecedented complexity. It will generate ~600 µs long *macropulses* of up to ~3000 bunches at a repetition rate of 10 Hz. Individual bunches out of these trains will then be distributed to several undulator beamlines and beam dumps using a system of fast intra-train kickers. Even the sheer number of subsystems and components—more than 1400 electromagnets for beam guidance and focusing, more than 450 beam position monitors, to pick two popular examples—makes it clear that a high degree of automatization will be needed to operate the machine efficiently. Control software should offer high level abstractions for all important machine parameters rather than forcing operators to work at the level of individual technical subsystems.

We have formed an interdisciplinary team of physicists, engineers, and computer scientists from multiple DESY groups with the goal of establishing the best possible operability for the machine. The scope of this *high level controls* (HLC) effort is a deliberately broad one. It includes the preparation of user friendly tools for various physical or technical tasks in the control room as well as the development of feedbacks and of middle layer servers. A lot of work is invested to improve the entire software stack of libraries, toolkits, and network services. This paper gives an overview of the various projects and activities, starting at the infrastructure level and closing with a discussion of stand-alone tools. For obvious reasons, we limit ourselves to the description of a few selected examples for each category.

[1] European X-Ray Free-Electron Laser Facility GmbH, Albert-Einstein-Ring 19, 22761 Hamburg, Germany
[2] Deutsches Elektronen-Synchrotron DESY, Notkestraße 85, 22607 Hamburg, Germany

PROTOCOLS AND INFRASTRUCTURE

Three main control system protocols are used at the XFEL: DOOCS [5, 6], Karabo [7], and TINE [8, 9]. Much work has been invested to improve the interoperability of these protocols [10], so that the problem of network communication across protocol boundaries is mostly solved. For HLC applications, efforts have concentrated on the development and improvement of interfaces for Matlab [11] and Python.

A *configuration database* has been set up as a central network service using DESY's Oracle 12c database system. It stores a complete list of beamline components and associated information such as calibration data. In this way, inconsistencies in the configuration of distributed servers can be minimized [12].

Finally, a *virtual accelerator* has been created by duplicating most of the middle layer and data acquisition software of the XFEL and exchanging the real front-end servers with dummies generating simulated data. This virtual XFEL has proven invaluable for the testing of high level software [13].

LIBRARIES AND TOOLBOXES

A multitude of libraries has been created to support the development of middle layer servers in C++. These projects include generic class libraries for access to various accelerator components, collections of algorithms for numerical tasks and fitting routines, classes for image analysis and conversion, an easy-to-use interface for retrieving information from the central configuration database, and various utility functions.

Matlab has traditionally been the main high-level interpreted language for physics applications at DESY. A broad assortment of utility functions is therefore available and actively maintained; a report on the DataGUI library for the design of advanced graphical user interfaces is found in [14]. We are currently focusing on the implementation of Matlab classes for access to specific accelerator components.

Recently, Python has been introduced as a second high-level language. Its main application at the moment is the rapid development of graphical user interfaces using the QT framework. Toolkit and library support is still fragmentary, but expanding continuously.

In general, we are encouraging the use of modern language standards where possible. For example, big parts of our C++

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

THE NEW CONTROL SOFTWARE FOR THE CERN NA62 BEAM VACUUM

S. Blanchard, F. Antoniotti, R. Ferreira, P. Gomes, A. Gutierrez, B. Jenninger, F. Mateo, H. Pereira
CERN, Geneva, Switzerland
L. Kopylov, S. Merker, IHEP, Protvino, Russia

Abstract

NA62 is a fixed target experiment to measure very rare decays of Kaons at CERN Super Proton Synchrotron accelerator. The NA62 experiment line comprises several large detectors installed inside a vacuum vessel with a length of 250 m and an internal diameter of up to 2.8 m.

The vacuum installation consists of 170 remote controlled pumps, valves and gauges. The operational specifications of NA62 require a complex vacuum control system: tight interaction between vacuum controllers and detector controllers, including pumping or venting vetoes, and detector start-stop interlocks; most of the valves are interlocked, including the large vacuum sector gate valves; the vacuum devices are driven by 20 logic processes.

The vacuum control system is based on commercial Programmable Logical Controllers (Siemens PLC: S7-300 series) and a Supervisory Control And Data Acquisition application (Siemens SCADA: WINCC OA). The control software is built upon the standard framework used in CERN accelerators vacuum, with some specific developments. We describe the control architecture, and report on the particular requirements and the solutions implemented.

INTRODUCTION

At CERN, the European Organization for Nuclear Research, the NA62 experiment analyses the outcomes of collisions of proton beams from the SPS (Super Proton Synchrotron) against a fixed target. The rare decays of Kaons are detected along a vacuum vessel of the 250 m. This large vacuum installation is broken down into 7 volumes, called vacuum sectors, separated either by sector valves or by windows. The use of sectors reduces mechanical intervention impact and sector valves, interlocked on pressure rise, avoid the propagation of leaks. The vacuum pressure specification for sectors 1 and 2 is lower than 10^{-2} mbar; it is achieved using 5 primary pumping groups. The vacuum pressure specification for sectors 3 to 7 is lower than 10^{-5} mbar; it is achieved using 8 turbo-molecular pumping groups and 7 large cryogenic pumping groups.

The control software for NA62 Beam Vacuum is based on the same framework as the vacuum control of the LHC (Large Hardon Collider) and its injectors [1].

The huge NA62 vacuum vessels require very large cryogenic pumping stations that offer very high pumping speed. In addition, usual turbo-molecular pumping is installed. The large detectors of NA62 are installed inside the beam vacuum vessels and require a complex hardware and software alarm system between vacuum and detector controls. Performant diagnostic tools are required to improve interventions and reduce costly beam interruptions.

New control device types and new control functionalities have been especially developed in 2014 for the NA62 Experiment. The software architecture is a two-layer system. The lower one, or automation layer, comprises a set of functions and data blocks, running in a Siemens Simatic S7™ PLC. The higher one, or supervisory layer (SCADA), was developed using the Siemens WINCC OA™ application; it consists of a set of routines and processes for monitoring, diagnostics and archiving, together with the graphical user interface.

DATABASE

The database structure, generically called Vac-DB, was developed in collaboration with IHEP for the vacuum control of the LHC and its injectors. This structure is also used for NA62.

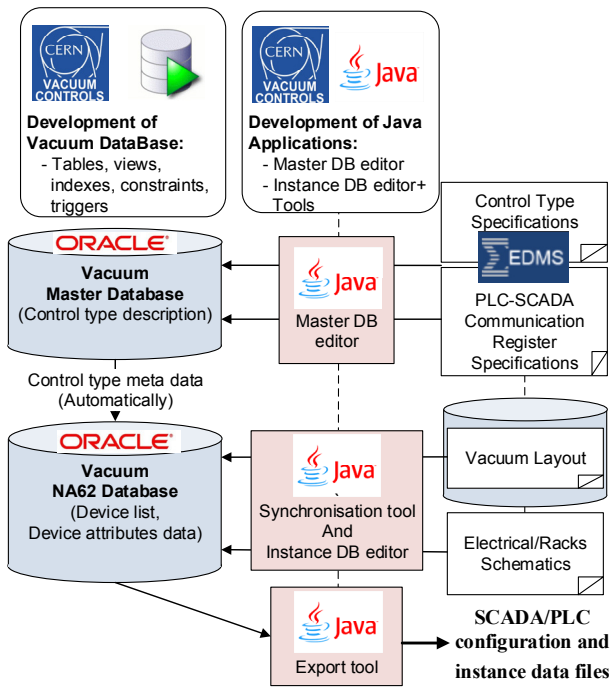


Figure 1: Vacuum database and applications.

Vacuum devices are classified into control types. The control type is identified by the family, type and sub-type of the device; it represents field devices (valves, mechanical pumps, active gauges, passive gauges, etc.) or virtual devices (pumping group process, gas injection process, software interlock, etc.).

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

THE UPGRADE OF CONTROL HARDWARE OF THE CERN NA62 BEAM VACUUM

F. Mateo, F. Antoniotti, S. Blanchard, R. Ferreira,
P. Gomes, A. Gutierrez, B. Jenninger, H. Pereira
CERN, Geneva, Switzerland

Abstract

NA62 is the follow-up of the NA48 experiment, in the SPS North Area of CERN, and reuses a large fraction of its detectors and beam line equipment. Still, there are many new vacuum devices in the beam line (including pumps, valves & gauges), which required a thorough modification of the control system and a large number of new controllers, many of which were custom-made.

The NA62 vacuum control system is based on the use of PLCs (Programmable Logic Controllers) and SCADA (Supervisory Control and Data Acquisition). The controllers and signal conditioning electronics are accessed from the PLC via a field bus (Profibus); optical fibre is used between surface racks and the underground gallery.

The control hardware was completely commissioned during 2014. The nominal pressure levels were attained in all sectors of the experiment. The remote control of all devices and the interlocks were successfully tested.

This paper summarizes the architecture of the vacuum control system of NA62, the types of instruments to control, the communication networks, the hardware alarms and the supervisory interface.

INTRODUCTION

NA62 will study the rare decay $K^+ \to \pi^+ \nu\bar{\nu}$ at the CERN SPS. NA62 is built to improve the current experimental precision by measuring kaon decays-in-flight with calorimetry to veto extra particles, very light mass trackers to reconstruct the momenta of the K^+ and the π^+ and full particle identification capability [1].

To reach the required sensitivity and acceptance, the vacuum in the decay area needed to be improved by at least two orders of magnitude compared to NA48. Many detectors that would normally be outside the vacuum chamber are now part of the vacuum system leading to a significant increase of the gas load to be pumped. A particularity of NA62 is the distribution of detectors along 250 m (Figure 1).

The NA62 experiment went through a complete refurbishment with the introduction of new detectors filed with gas. Their operation requires a vacuum system, based on turbo molecular and cryogenic pumps. The beam line is divided into 7 sectors, each having specific vacuum requirements. The vacuum equipment in each sector will be described in the next section.

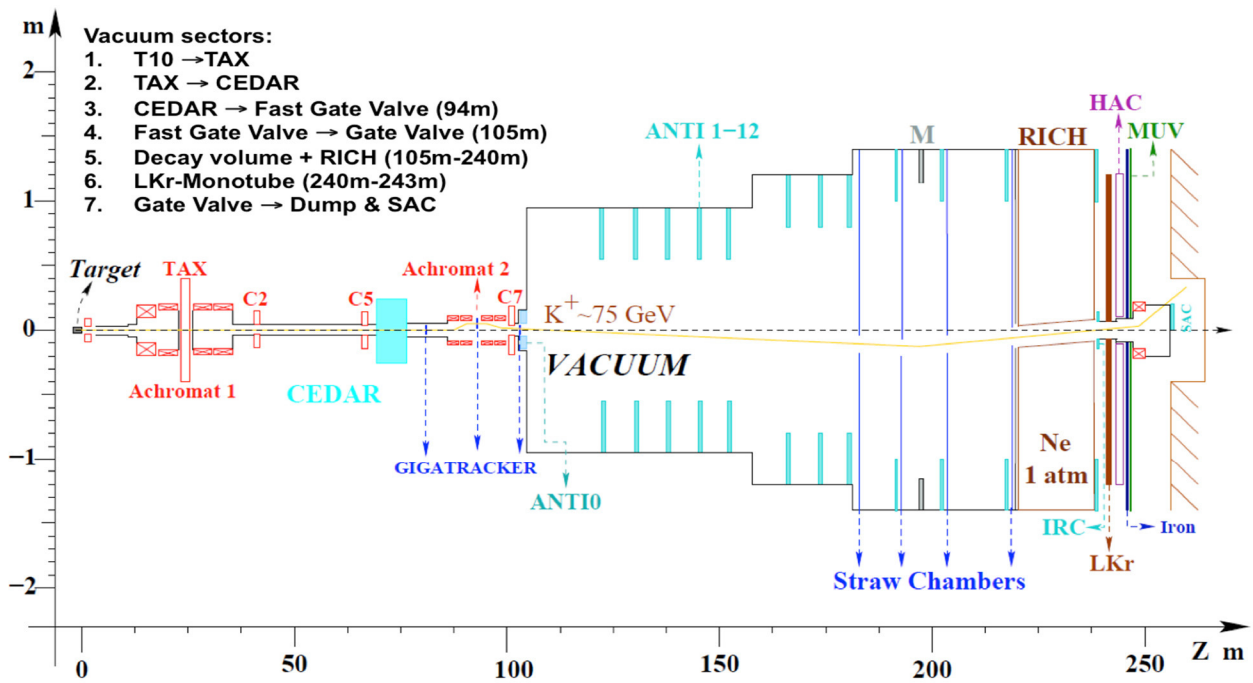


Figure 1: Layout of the NA62 experiment

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

CONSOLIDATIONS ON THE VACUUM CONTROLS OF THE CERN ACCELERATORS, DURING THE FIRST LONG SHUTDOWN OF THE LHC

P. Gomes, F. Antoniotti, F. Aragon, F. Bellorini, S. Blanchard, J-P. Boivin, N. Chatzigeorgiou, F. Daligault, R. Ferreira, J. Fraga, J. Gama, A. Gutierrez, P. Krakowski, H. Pereira, G. Pigny, P. Prieto, B. Rio, A. Rocha, H. Vestergard
CERN, Geneva, Switzerland
L. Kopylov, S. Merker, M. Mikheev ; IHEP, Protvino, Russia

Abstract

For two years (Spring 2013 – Spring 2015), the LHC went through its first long shutdown (LS1) [1]. It was mainly motivated by the consolidation of magnet interconnects, to allow operation with 6.5 TeV proton beams. Moreover, around the accelerator complex, many other systems were repaired, consolidated or upgraded, and several new installations came to life.

The standardization of vacuum controls has progressed in the injectors, with the renovation of most of their obsolete equipment.

In the LHC, many new instruments were added, the signal transmission integrity was improved, and the exposure to radiation was reduced in critical places. Several developments were needed for new equipment types or new operational requirements.

INJECTORS & LHC CONSOLIDATIONS

Complex PS

The LS1 was an opportunity to extend the standardization of vacuum controls to the Proton Synchrotron ring (PS), the Antiproton Decelerator (AD), and the nTOF experiment. The legacy controls were updated to the same PLC/SCADA-based architecture as in the other machines. Only the pumping groups of the PS complex remain to be upgraded, during the following years. [2]

The existing controllers were either upgraded, or replaced by newer generations, connecting to the PLC directly through Profibus or via remote-IO stations. This involved 500 cables and 700 instruments.

Above all, a large effort was put in the production of complete and up-to-date documentation on the control system layout and on its components technical details.

Linac2, at the head of the proton acceleration chain, received a new set of six pumping groups. These were especially developed to compensate for the dying old ion pumps, helping it to survive until replaced by Linac4.

SPS

In the SPS, controls renovation was applied to its injection line, the interfaces for the Beam Interlock System, and COLDEX.

Several new gauges and valves were added, for a few new sectors. Furthermore, all racks and cables were prepared for the new re-sectorization of all the ARCs.

Campaigns for replacement of aged cables were performed in several places. In total, some 300 cables were replaced, and 200 other were newly installed, mainly for some 80 new instruments.

LHC

About 600 cables and 400 gauges, valves, and pumps, were newly installed in the LHC, together with the corresponding interlocks. Some of them required specific controls developments, treated in detail further down.

The signal integrity has been revised and improved, for the gauges in the ARCs and for the bypass-valves. To mitigate radiation effects in Point7, 27 racks full of equipment were dismounted and relocated into a less exposed zone, while 400 cables had to be extended. These subjects are treated in another paper [3].

During the second half of 2014, six teams of two people were full-time involved in the commissioning of all controls and instrumentation of the LHC. All instrument chains were fully tested; and calibrated when necessary. Documentation was updated about racks layout and interlocks cabling.

During the LHC global check-out, systematic tests were performed on all interlocks, to every sector valve, and to the beam interlock system, automatically from SCADA.

All in all

In total, more than 2 900 cables and 1 500 instruments were commissioned in the accelerators consolidations and new projects. The controls team had to be reinforced by 50%, mainly with engineers from collaborating Institutes, trainees, and technicians from industrial support. The main challenges were the coordination of field activities with access restrictions and other teams' tasks, and having each machine commissioned and ready in time for beam.

DEVELOPMENTS FOR PUMPS

NEG Pumps

A Non-Evaporable Getter (NEG) is a highly reactive and porous material; it needs to be activated, by heating under vacuum, with a suitable temperature and duration. The recommended activation temperature is around 475 °C, corresponding to 6A under 16V, for 1 h.

To increase pumping speed in critical equipment in the LSSs of the LHC, NEG cartridges (Capacitorr D400 2 from

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

DEVICE CONTROL DATABASE TOOL (DCDB)

Pavel Maslov, Matej Komel, Miroslav Pavleski, Klemen Zagar, Cosylab, Ljubljana, Slovenia

Abstract

We have developed a control system configuration tool, which provides an easy-to-use interface for quick configuration of the entire facility. It uses Microsoft Excel as the front-end application and allows the user to quickly generate and deploy IOC configuration (EPICS start-up scripts, alarms and archive configuration) onto IOCs; start, stop and restart IOCs, alarm servers and archive engines, and more. The DCDB tool utilizes a relational database, which stores information about all the elements of the accelerator. The communication between the client, database and IOCs is realized by a REST server written in Python. The key feature of the DCDB tool is that the user does not need to recompile the source code. It is achieved by using a dynamic library loader [1], which automatically loads and links device support libraries. The DCDB tool is compliant with CODAC (used at ITER and ELI-NP), but can also be used in any other EPICS environment (e.g. it has been customized to work at ESS).

INTRODUCTION

In a physics facility containing numerous instruments, it is advantageous to reduce the effort and repetitiveness needed for changing the control system (CS) configuration: adding new devices, moving instruments from beamline to beamline, etc. We have developed the Device Control Database (DCDB) tool [2], which provides an easy-to-use interface for quick configuration of the control system for the entire facility. The DCDB-tool allows the user to quickly generate and deploy configuration for input/output controllers (IOCs) (EPICS start-up scripts, alarms and archive configuration) onto IOCs; start, stop and restart IOCs, alarm servers and archive engines, and more.

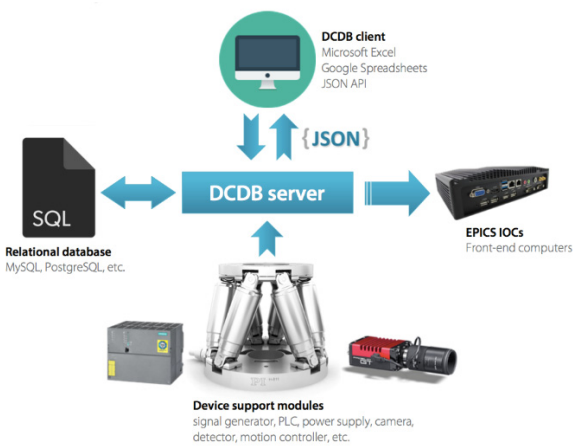


Figure 1: DCDB architecture.

DCDB ARCHITECTURE

The DCDB-tool uses a MySQL relational database. The backend is a typical web-server (Fig. 1), which is realized with a combination of the following Python modules: flask-restful (REST server), sqlalchemy and pymysql (database communication layer), and paramiko (ssh). The front-end is a Microsoft Excel plugin written in C# using .NET technology. IOCs are Linux machines running EPICS and procServ [3]. The client-server communication is based on the exchange of JSON objects (strings).

DEVICE SUPPORT MODULES

To start using the DCDB-tool, simply prepare EPICS device support modules and register them with the DCDB server. In general, device support modules are created in the standard EPICS way. The DCDB-tool introduces the idea of using dynamic macros in the startup scripts. Three additional files, namely: init.cmd, init-pre.cmd and init-post.cmd (Fig. 2) contain macro definitions to be stored in MySQL, and IOC shell commands that register and setup the support module before and after IOC initialization.

```
.
├── db
│   └── BeamPositionMonitor.db
├── dbd
│   └── BeamPositionMonitor.dbd
├── init.cmd
├── init-post.cmd
├── init-pre.cmd
└── lib
    └── linux-x86_64
        ├── libBeamPositionMonitor.a
        └── libBeamPositionMonitor.so
```

Figure 2: Files to deploy.

The procedure of creating device support modules for ITER using special extensions to the *mvn iter* plugin was described in [4]. In the latest version we have tailored DCDB to work in the new EPICS environment developed at ESS [5], since they have a similar approach to using dynamic library loading, and in addition to our setup support dependency resolution and have a simpler unit development workflow. The latter is achieved by creating a Makefile, adding source files and running make. If you want to locally install the module, run the shell command: ***sudo make install***. In order to register the module with the DCDB tool, run ***make dcdb.import***. To delete the module from DCDB, run ***make dcdb.delete***.

The last two commands require two environmental variables *$(DCDB_SERVER)* and *$(DCDB_PORT)* to correspond with the parameters of a running DCDB server.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

DESIGN STRATEGIES IN THE DEVELOPMENT OF THE ITALIAN SINGLE-DISH CONTROL SYSTEM

A. Orlati, M. Bartolini, S. Righini, INAF-IRA, Bologna, Italy
M. Buttu, A. Fara, C. Migoni, S. Poppi, INAF-OAC, Cagliari, Italy

Abstract

The Italian National Institute for Astrophysics (INAF) manages three radio telescopes: the Medicina and Noto dishes and the newly-built SRT. In order to make their capabilities more valuable to the scientific community, we started the DISCOS (Development of the Italian Single-dish COntrol System) project. DISCOS is implemented according to a distributed Component-Container model and hides to the users the differences among the telescopes by presenting the same user interface and the same data format. The complexity of coping with three heterogeneous instruments was handled designing a software development infrastructure with a wide monolithic codebase (libraries, components and generic interfaces), which is completely shared among the three product lines. This design permits to produce new software components with a minimum effort and to set up the same test suites for all the environments, thus leading to an affordable development and maintenance process. In this paper we illustrate the design strategies and the development techniques used to realize and optimize this common control software. We also provide a description of the project status and future plans.

INTRODUCTION

The National Institute for Astrophysics (INAF) manages three radio telescopes in Italy: The Sardinia Radio Telescope (SRT) [1] and the Medicina and Noto 32-m dishes. The newly-built SRT, located in the Sardinia island, was inaugurated in 2013. Since the early stages of the project, staff from all the three telescopes has been involved in the development of the SRT control software. This forced us to cope with both the development of a brand-new system and the maintenance of the already existing telescopes, compelling the involved personnel to learn and enhance their competence in distinct and heterogeneous systems based on completely different technologies. Developing a control software to be installed at all the telescopes was the natural and straightforward approach to this problem.

This idea was formalized with the creation of the Development of Italian Single-dish COntrol System (DISCOS). The project aims to provide all three Italian radiotelescopes with a common infrastructure that increases their capabilities and the usability for the scientific community and, at the same time, optimizing the efforts made by the technical staff for its development and maintenance.

DISCOS is based on ALMA Common Software (ACS) developed at ESO for the ALMA project [2]. ACS implements a Component-Container model via CORBA (Common Object Request Broker Architecture). It provides a set of tools, libraries and development patterns that hide the complexity of CORBA. This approach permits to reduce the time required for coding and development. The framework also supports both real and non-real time platforms for C++, Python and Java.

The core of the project was completed during the main development stage. The SRT and Medicina installations are now fully operational and supervise all the telescope operations. The Noto antenna is also equipped with a preliminary version, which is going to be finalized in a few months.

In this paper we describe the salient design choices that allowed us to reuse the most part of the code for all the production lines and to hide the instrument complexity under a common infrastructure. We also illustrate the maintenance and development workflows, as they are fundamental part of the DISCOS design.

THE TELESCOPES

The INAF radiotelescopes share some common aspects but are in general different from one another. Medicina and Noto have passed through years of development and updates, while SRT is the result of an organic employment of state-of-the-art technologies. Table 1 recaps the main characteristics and the differences, listing all the major sections – servo systems, front-ends, back-ends. All the telescopes are presently participating in the VLBI network. Medicina and Noto are also offered to the scientific community as single-dish facilities (SRT will very soon be, as well).

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

TANGO INTEGRATION OF A SPECIFIC HARDWARE THROUGH HTTP-SERVER

A. Panov, A. Korepanov

The Budker Institute of Nuclear Physics, Novosibirsk, Russia

Abstract

MAX IV are new synchrotron third generation. MAX IV synchrotron consist of 1.5 GeV storage ring, 3.0 GeV storage ring and linac; it is located in Lund, Sweden. Structure of storage rings contains several pulse magnets (kicker and pinger). Control system of pulse power supplies based on LTR crate with several modules (ADC, DAC, input/output registers etc.). LTR crate is a product Russian firm L-CARD. In order to communicate with crate native LTR-server is used. LTR-server is a Windows application based on use of sockets. Control system of MAX IV and Solaris uses TANGO. For integration LTR-crates in final structure, special software gateway is used. This gateway is a set of several specific Windows applications implemented by using Qt5 libraries. Gateway allows communicating TANGO-server with crate through built-in HTTP-server.

SPECIFIC HARDWARE

The MAX IV facility included two storage rings for the production of synchrotron radiation: the 3 GeV and 1.5 GeV storage rings [1]. Both rings will be operated with top-up shots supplied by the 3.5 GeV MAX IV linac acting as a full-energy injector at up to 10 Hz repetition rate. At the commissioning phase, the single dipole kickers will be used for the injection in the both rings. In addition, pulsed vertical kickers (pingers) are needed for performing tests with the stored electron beam.

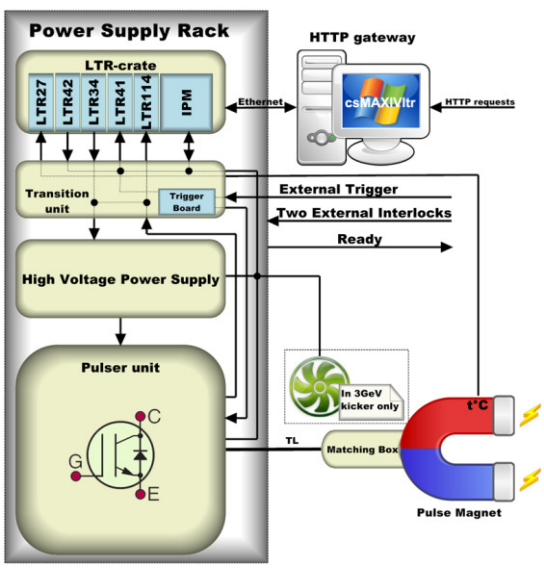


Figure 1: General pulse power supply structure.

Table 1: Pulse Magnet Systems Specifications

	1.5GeV		3GeV	
	kicker	pinger	kicker	pinger
Integrated field, Gm	143	75	450	110
Peak current, A	1060	2050	3600	2600
Pulse length, μs	0.64	0.64	3.5	3.5
Maximum rep. rate, Hz	10	1	10	1
Pulse to pulse timing,ns	±5	±5	±5	±5
Pulse to pulse repeatability, %	±0.1	±1	±0.1	±1

The requirements for the pulse magnet systems are listed in Table 1. Structure of all system (power supply with control electronics and magnet) are illustrated on Fig.1.

Control system of the pulse power supply based on 19" LTR-crate (LTR-EU-16). It is a product of a Russian firm L-CARD (http://lcard.ru). Base structure of the LTR-crate are represented on a Fig. 2.

This crate include FPGA Cyclon EP1C3 for

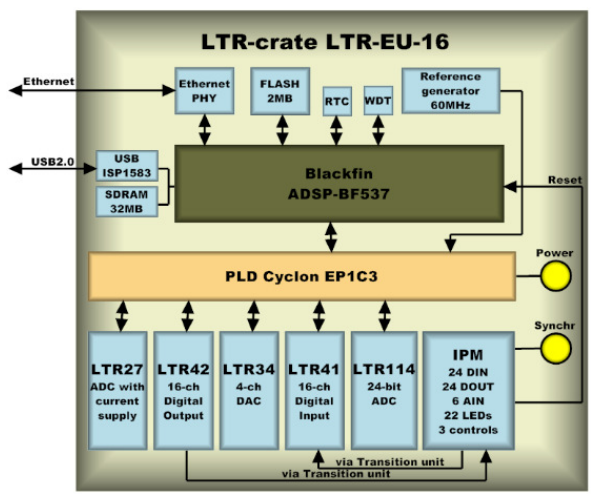


Figure 2: Power supply control system (LTR crate with modules) block diagram.

communication with installed in crate modules via proprietary interface and digital signal processor Blackfin ADSP-BF-537 for preliminary data processing. DSP operates under control of Visual DSP++ Kernel (VDK) [2]. Sources of a DSP firmware in free access and can be modified for individual requirements. Crate can communicate with PC over USB2.0 or Ethernet (TCP/IP) via L-CARD specific protocol.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

MEASUREMENTS, ALARMS AND INTERLOCKS IN THE VACUUM CONTROL SYSTEM OF THE LHC

G. Pigny, F. Antoniotti, J.P. Boivin, N. Chatzigeorgiou, J. Gama, P. Gomes, P. Krakowski, H. Pereira, H. Vestergard
CERN, Geneva, Switzerland

Abstract

In the LHC beam pipes and cryostats, the pressure measurement covers a wide range, from 1500 mbar down to 10^{-11} mbar and even lower. If vacuum deteriorates, alarm signals are generated and sent to other systems, like cryogenics, accelerating cavities, kicker magnets. In addition, an unacceptable pressure rise in beam pipes generates interlocks to close the adjacent sector valves, thus isolating the sector, so that the pressure rise does not propagate. This pressure interlock is simultaneously forwarded to the beam interlock system.

This paper describes the instrumentation, the interlocks and alarms logic used in the vacuum control system of the LHC. We analyze the possible signal degradation caused by ionizing radiation or due to cable length, shielding and grounding. During the first LHC long shut down, several corrections were applied to mitigate radiation effects and improve signal integrity. The tests performed for the vacuum control commissioning and LHC machine checkout are also presented.

VACUUM MEASUREMENT

Vacuum gauges in accelerators have to cover about 15 decades of pressure. The direct pressure gauges, which measure the force per unit area, are gas-type independent. They cover the rough and medium vacuum ranges. The indirect pressure gauges, which measure a pressure-dependent property of the gas, are gas-type dependent. They can cover a wider range of vacuum, from the rough to the ultra-high vacuum range [1].

The membrane gauge covers the range from 0 to 2 000 mbar. It is composed by a thin silicon crystal layer used as a membrane, whose deformation induces a piezo-resistive effect on the four resistances of a Wheatstone bridge. The resistance variation is a function of the pressure. The signal to be read across the bridge is in the mV range.

The Pirani gauge covers the range from 1 000 to 10^{-4} mbar. The measurement is based on the dependence of the thermal conductivity on the gas pressure. A filament is heated and maintained at a constant temperature. The current needed to maintain this temperature depends on the gas conductivity and therefore is a measure of the pressure. The signal delivered by the amplifier is in the Volt range. Above 1 mbar the measurement becomes non-linear, and depends on the gas species

The Penning gauge covers the range from 10^{-5} down to 10^{-11} mbar; it is also called inverted magnetron cold cathode gauge. It is composed by a thin cylindrical anode rod, surrounded by a cathode. A high voltage (HV) potential of 3kV creates a high electric field between the two electrodes, which favors the ionization of gas molecules. Around the electrodes, a permanent magnet produces a field that forces the free electrons into spiral paths, thus increasing the probability of ionizing more molecules. The ionization current is a function of the pressure. In the low range of measurement, the current can be as low as the pA; a high quality HV triaxial cable is used for signal transmission.

In the range from 10^{-5} down to 10^{-9} mbar, the Sputter Ion Pump can be used to measure the pressure, in addition to its pumping function. It is composed by several Penning cells, with the particularity of Titanium cathodes and a potential of 7 kV. The principle of vacuum measurement is the same as for the Penning gauge. In the low range of measurement, the current can be as low as 100 nA, and high quality HV triaxial cable is used for signal transmission.

Finally, to measure down to 10^{-12} mbar, the Bayard-Alpert gauge is used. A heated filament emits electrons which are accelerated by the grid potential (+150 V). On their path, these electrons ionize gas molecules, which are gathered by a collector. The collector current is a measure of the ionization and therefore of the gas pressure. It can be as low as 100 fA. High quality triaxial cable is used for this measurement.

INTERLOCKS AND ALARMS

Alarms are binary signals sent to the operators' console, e-mail or SMS, to draw their attention to particular equipment status or situations. There are also alarms directly cabled to other systems, for information or to be used in their control logic.

Interlocks are binary signals used within the vacuum system to prevent entering undesired states. They have a higher priority than the normal process logic or the operator's commands. Apart from the beam permit interlock, the vacuum control system does not interlock any other external systems directly.

The detection of a pressure rise, above a predefined threshold, can be used to produce alarms or interlocks. The gauges used for alarms or interlocks can be the Membrane, Pirani, Penning and Ion Pump [2].

Interlocks for the Vacuum Sector Valves

The beam pipes are divided in "vacuum sectors" that can be isolated by "sector valves" (VVS), to avoid the propagation of leaks over a large volume.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

CONTROLS AND INTERLOCKS FOR THE NEW ELETTRA SUPER CONDUCTING WIGGLER

L. Pivetta, F. Giacuzzo, G. Scalamera, D. Vittor,
Elettra Sincrotrone Trieste, Trieste, Italy

Abstract

During the last two years, triggered by the construction of the XRD2 beamline, and to comply with the topup operations, a complete refurbishment of the Elettra Super Conducting Wiggler (SCW) has been carried out. Alongside with the mechanical, cryogenic and electrical components, also the electronics, the control and interlock systems have been upgraded. The MVME5110 PowerPC single board computer, which is a standard in the Elettra control system, has been adopted, as well as RS232 communication modules, analog to digital converters and digital I/O lines. To cope with the high output power of the SCW, up to 18 KW, the interlock system, protecting both the wiggler and the beamline frontend, has been completely redesigned. The control system software has been rewritten from scratch using the TANGO software framework. The complete system has been tested during the second half of 2014 and is now fully operational.

INTRODUCTION

A 3.5 T superconducting wiggler, designed and built by the Budker Institute of Nuclear Physics (BINP), has been installed in the Elettra storage ring as a photon source for the X-ray diffraction beamline XRD2 in 2002 [1, 2]. Due to the lack of funding for the beamline construction, the SCW didn't came in operation in the subsequent years. Then, in 2011, a collaboration agreement between the Indian government and Elettra Sincrotrone Trieste provided the funds required for the beamline construction and the SCW refurbishment. During 2012 and the first months of 2013, at the BINP laboratory in Novosibirsk, the SCW has been provided with a new cryostat aimed at decreasing the helium consumption and increasing the reliability. Then, in July 2013, after successful site acceptance tests, the SCW has been installed in the Elettra storage ring [3]. A picture of the installed wiggler is shown in Fig. 1.

Simplifying, the main components of the SCW are the chassis, the cryostat, the liner and the superconducting coils; the principle of operation of the SCW consists in keeping the cryostat temperature near to the 4.5 K, using liquid helium, and feeding the superconducting coils with the appropriate currents in order to reach the desired magnetic field.

ELECTRONICS

The original Danfysik MPS883 power supplies, powering the central and the side superconducting coils, after being revised have been kept in operation, as well as the four existing Leybold Coolpack 6000 compressors, provided with new Coolpower cooling heads. A new junction box, shown in Fig. 2, has been designed by BINP to replace both the old cabled box and the VME-based signal conditioning and acquisition board. The new junction box, based on an Atmel ATmega128 microprocessor, acquires the signals coming from the temperature and pressure sensors installed into the cryostat and implements the quench interlock logics. If a quench occurs the junction box turns off the power supplies that feed the superconducting coils driving their external interlock input. A RS232 serial line is available for remote control and supervision.

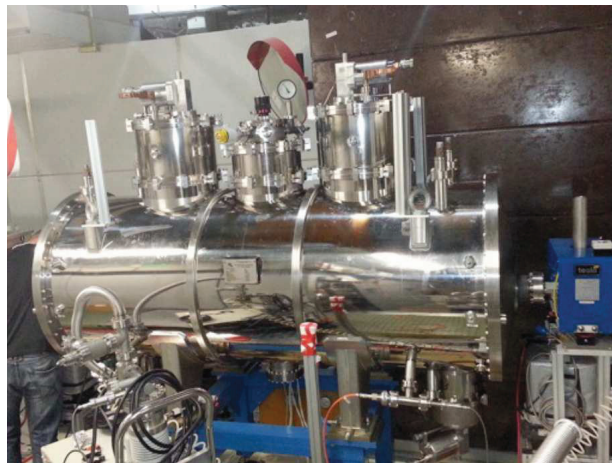


Figure 1: The refurbished super conducting wiggler installed in the Elettra storage ring.



Figure 2: The new Junction Box.

Ion pumps, powered by the DUAL high voltage power supplies manufactured by Varian, are also installed to provide the requested vacuum levels. A pirani vacuum gauge allows to monitor the pressure in the vessel insulation, whilst some penning gauges are used in the front-end vacuum chamber; they are all powered with two TPG300 power supplies.

CONTROLS

In conjunction with the refurbishment of the cryostat the SCW control system has also been updated. The legacy control system, based on a VME crate and a Motorola 68K microprocessor single board computer running VxWorks, provided by BINP as a turnkey system, has been dismantled. Also, the ELTEC BAB40 single board computer, run-

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

CONTROLS INTERFACE INTO THE LOW-LEVEL RF SYSTEM IN THE ARIEL e-LINAC AT TRIUMF

J.J. Pon, K. Ezawa, R. Keitel, R.B. Nussbaumer, J.E. Richards, M.A. Rowe, P.J. Yogendran
TRIUMF, Vancouver, BC, Canada

Abstract

Phase I of TRIUMF Advanced Rare Isotope Laboratory (ARIEL) was completed in September 2014. At phase I, the Low-Level RF (LLRF) system of ARIEL's electron linear accelerator (e-Linac) consists of a buncher and a deflector, one single-cavity injector cryomodule and the first cavity of two dual-cavity accelerating cryomodules. The model for the e-Linac LLRF system is largely based on the experience gained from the fully-commissioned TRIUMF ISAC-II linear accelerator (linac). Similarly, the EPICS-based Controls for the e-Linac LLRF builds on the lessons learned with the linac LLRF Controls. This paper describes the interface between the ARIEL Control System (ACS) and the e-Linac LLRF using EPICS ASYN/StreamDevice and a SCPI-like protocol. Also discussed are the ACS EDM displays and future plans for LLRF Controls.

INTRODUCTION

The Advanced Rare Isotope Laboratory (ARIEL) is the latest project in the Rare Isotope Beam (RIB) program at TRIUMF. In conjunction with ISAC and ISAC-II, ARIEL is designed to triple RIB production and expand the range of exotic isotopes “for Nuclear Physics and Astrophysics, Nuclear Medicine and Materials Science” [1].

In September 2014, phase I of ARIEL was completed with the delivery of a 300KeV DC thermionic gun, a buncher and a deflector, a single-cavity injector cryomodule (ICM), and the first cavity of a dual-cavity accelerating cryomodule (ACM). These components constitute the first half of ARIEL's superconducting electron linear accelerator (e-Linac). The second half of the e-Linac is presently under way and comprises the second cavity of the first ACM, and a second dual-cavity ACM, for a total of five cavities in three superconducting cryomodules.

The design of the Low-Level RF (LLRF) system for ARIEL's e-Linac borrowed extensively from the model used in ISAC-II superconducting linear accelerator (linac) [2] – TRIUMF's first fully-commissioned linac. It was natural then, that the new ARIEL Control System (ACS) interface into the e-Linac LLRF also followed the model used in ISAC-II linac Controls. However, as in any new project, LLRF Controls of ARIEL's e-Linac presented an opportunity to improve on some of the unexpected drawbacks encountered in ISAC-II Controls.

INTERFACE

LLRF-ACS Interface

The overall framework used in ISAC-II LLRF Controls has been carried over into the LLRF-ACS interface. At one end, the LLRF system exchanges information with Controls via an EPICS IOC. At the other end, the IOC makes LLRF data available to EPICS clients in the form of process variables (PV). While this model has worked well at TRIUMF, the way it was implemented in ISAC-II LLRF Controls uncovered some drawbacks that were addressed in the LLRF-ACS interface.

In ISAC-II, the LLRF system interfaced with the Controls EPICS IOC using a shared memory model [3]. The implementation of this method brought in a number of firm constraints. First, inter-process communication via shared memory implies a high degree of coupling. The LLRF system and ISAC-II Controls had to adhere to a strict memory mapping scheme, was not very portable, and carried a moderate amount of complexity to maintain.

Another implication on the use of shared memory is that the IOC needs to run in the same LLRF system computer. The host machine is owned and maintained by the LLRF group. For simplicity and security reasons, LLRF computers are stand-alone machines and are not connected to the site's network. If Controls wanted to perform diagnosis or maintenance of the IOC it was necessary to inconvenience the LLRF group for physical access to the host computer. Even simple actions like viewing IOC log messages or machine crash dumps required local access to the machine.

Furthermore, the IOC had to be built to run on Microsoft Windows. But Controls have had limited experience developing for the Windows environment. Only one Controls programmer gained enough knowledge to build a heavily-customized Windows IOC based on the EPICS Portable Channel Access Server framework and make available to LLRF engineers a C++ API to read/write the shared memory. Plus, the programming was done in Visual Studio and involved the use of WIN32 DLLs [4], these development tools were also unfamiliar to Controls.

During the early stages of ARIEL phase I an alternative implementation for the LLRF-ACS interface was put forward. After some consultation with the relevant groups one of the main conclusions was to not proceed with the shared memory model. Instead, the ACS would use the EPICS ASYN/StreamDevice method of communication with LLRF. The Controls group is familiar with

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

LABVIEW AS A NEW SUPERVISION SOLUTION FOR INDUSTRIAL CONTROL SYSTEMS

O.O Andreassen, F. Augrandjean, E. Blanco Vinuela, D. Abalo Miron, M. F. Gomez De La Cruz, A. Rijllart, CERN, Geneva, Switzerland

Abstract

To shorten the development time of industrial control applications, CERN has developed the Unified Industrial Control System (UNICOS) framework, which standardizes the programming of the front-ends and the configuration of the Supervisory Control and Data Acquisition (SCADA) layer. At CERN the SCADA system of choice is WinCC OA, but for some specific projects (small or initial prototypes not connected to accelerator operation or not located at CERN) a more customisable supervision using LabVIEW is an attractive alternative. Therefore, a system called UNICOS in LabVIEW (UiL), has been implemented. It provides a set of highly customisable re-usable components, devices and utilities. Because LabVIEW uses different programming methods than WinCC OA, the tools for automatic instantiation of devices on the supervision layer had to be re-developed, but the configuration files of the devices can be reused. This paper reports how the implementation was done, it describes the first project implemented in UiL and an outlook to other possible applications.

INTRODUCTION

The CERN Unified Industrial Control System (UNICOS) project proposes a standardized architecture for both, supervision and control layers. It also defines a set of tools for automatic instantiation of devices, in the two mentioned layers, employing re-usable components, devices and utilities.

However, for smaller lab-size solutions e.g. systems which are only expected to exist for a short time or when users do not have the time and resources to invest in a complete SCADA as WinCC Open Architecture (WinCC OA), a LabVIEW based solution would fill this gap [1][2].

UNICOS in LabVIEW (UiL) gives the user much of the same look & feel and functionality of WinCC OA as Human Machine Interface (HMI).

UiL generates automatically the supervision objects from the same configuration file provided by UNICOS Application Builder (UAB), which was developed to generate the code and configuration files of the different components of the UNICOS project (e.g. WinCC OA, Siemens PLCs etc.) [3][4]. These objects are the counterpart instances of the control objects deployed in the PLCs in the control layer.

Once the objects are instantiated in the UiL application, they can be used in the application synoptic by means of widgets, a reduced but still meaningful vision of its status, and faceplates, a full view of the object status and the place where the operator will interact with the control system.

In addition to the widgets and faceplates, UiL provides a trending tool, an alarm and an event list. All the components, widgets and user interaction in UiL are designed to resemble the UNICOS look and feel as much as possible; making sure that the transition from one SCADA tool to the other is effortless.

UNICOS

UNICOS was initially developed in 1998 as a need to develop the LHC cryogenics control system, but became a CERN de facto standard for industrial control systems. As shown in the figure 1, the framework connects, instantiates and interacts with both the supervision and control layers with commercial industrial components [1][5].

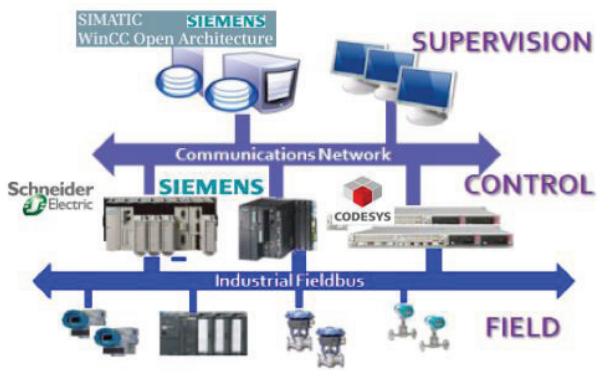


Figure 1: UNICOS architecture.

Currently UNICOS uses WinCC OA as a commercial SCADA system and both Siemens, Schneider, as Programmable Logic Controllers (PLC) [1][6]. A third possibility using Codesys extends the usage to industrial computers (IPC)

The possibilities and the application of the UNICOS framework extend to monitoring and supervision applications (e.g. Magnets alignment, Collimators, Quench protection systems...) where the control layer could be not based on UNICOS.

The UNICOS framework is based on an object oriented concept consisting of standard object classes featuring process control objects, controllers, valves, pumps, heaters, inputs, outputs etc. It also establishes how these objects interact and formalizes the control unit definition together with the way the specific control logic is implemented.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

THE CONTROL SYSTEM FOR TRIM-COIL RELAY-SELECTORS IN J-PARC MR

K. C. Sato*, S. Igarashi, N. Kamikubota, N. Yamamoto,
J-PARC, KEK & JAEA, Ibaraki-ken, Japan
S. Yoshida, Kanto Information Service (KIS), Accelerator Group, Ibaraki, Japan

Abstract

In J-PARC main ring, each of the main magnets (Bending, Quadrupole and Sextupole) has a trim-coil. The primary aim of trim-coil is to correct small deviation of each magnetic field. In addition, we use them for other purposes such as: (1) In Beam-Based-Alignment studies, (2) as flux monitors, and (3) to make a short-circuit to reduce ripples of magnetic field. At a moment, each trim-coil can be used for only one purpose. We developed a relay-selector system which enables selection of connection to equipment depending on purpose. When we switch the connection, we have to change 1,200 on-site relays manually, distributed in three buildings. Thus, a control system for trim-coil relay-selectors have been developed in 2014-2015. EPICS tools and environment are used to develop the system. The system comprises PLC I/O modules with a controller running EPICS on Linux. Its operation begun in April, 2015. By using the system, we expect much easier switching of relay-selectors and reduce overheads than before.

INTRODUCTION

J-PARC (Japan Proton Accelerator Research Complex) is a high-intensity proton accelerator facility. It has been operated collaboratively by Japan Atomic Energy Agency (JAEA) and High Energy Accelerator Research Organization (KEK). It consists of three accelerators: a linear accelerator (LINAC), a Rapid Cycling Synchrotron (RCS), and a Main Ring (MR). MR started beam operation in 2008 [1]. MR is an accelerator with circumstance of about 1570 meter. This paper aim to main magnets accelerator components.

There are 3 types in the MR main magnets: bending magnet (BM), quadrupole magnet (QM) and sextupole magnet (SM). The number of BM (QM, SM) is 96 (216, 72). They are arranged with a 3-fold symmetry in the MR tunnel. The main power supplies are installed in the MR power-supply buildings. As an example, the cabling layout of a typical family "QFN" is shown in figure 1. The power supply for QFN is located in the second power-supply building, and QFN magnets are wired clockwise. Each of these magnets has an independent trim-coil.

The trim-coil can be connected to four different equipment such as:

- "Harmonic" - to correct small deviation of each magnetic field.
- "BBA" - in Beam-Based-Alignment studies.

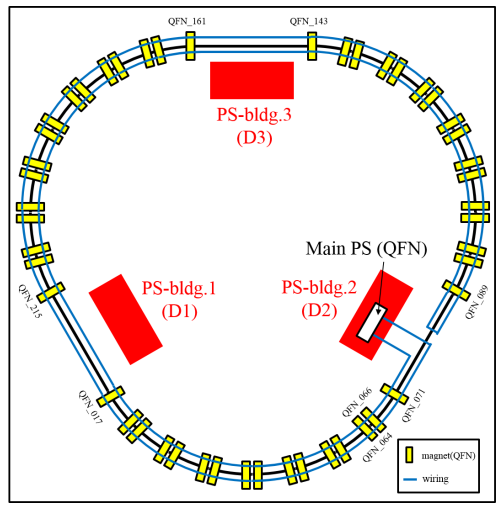


Figure 1: The layout of QFN power supply cabling.

- "Flux" - as magnetic flux monitors.
- "Short" - to make a short-circuit to reduce ripples of magnetic field [2].

In 2014-2015, we have developed the relay-selector system. Before the relay-selector system is installed, we have to switch connections between a trim-coil and a corresponding equipment by hands. All trim-coil connectors are placed in MR three power-supply buildings. So we need a lot of time to change connections. This relay-selectors system can reduce these overhead time since need not to switch the connection by hands.

By the way, the relay-selector system is applies to trim-coils of QM and SM. Since the BM trim-coils are used only for the "Short" purpose, it is independent from the system.

TRIM-COIL RELAY-SERECTORS SYSTEM

Each trim-coil need 4 relays. As shown in Figure 2 and 3, a relay-unit has 8 relays for 2 trim-coils. An aggregation-unit gathers cables from 8 relay-units. A control panel with buttons and fault indications is mounted in a rack. A relay-selector system, which consists of three rack, has 6 aggregation-units. Thus, a relay-selector system has to control 384 relays. J-PARC MR have three power-supply buildings, and each building has one relay-selector system. In total, this control system need to switch 1152 relays (equal to 288 trim-coils x4).

* kenichi.sato@j-parc.jp

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

DESIGN AND DEVELOPMENT OF THE ECR ION SOURCE CONTROL SYSTEM

Hyungjoo Son, Sangil Lee, Chang Wook Son, Hyojae Jang, IBS, Daejeon, S.Korea

Abstract

The Rare Isotope Science Project at the Institute for Basic Science constructs a heavy ion accelerator (RAON) facility in South Korea. The stable ion beam for the RAON accelerator could be generated by ECR ion source system. Therefore, it is necessary to build an ECR ion source control system that could be integrated into an accelerator control system easily. The vacuum control system is divided several parts because of one vacuum chamber among three different voltage stages (ground, 50 kV, and 80 kV). In this report, we will present the preliminary design and implementation of vacuum control system for the ECR ion source. We plan to use a Programmable Logic Controller (PLC) in order to control the vacuum system through interlock logic program. The PLC system has two major components: a digital I/O module that provides power to each component and standard RS-232 modules to connect the gauge and pump controllers. In addition, we will discuss its extension plan to integrate the vacuum control system into the RAON accelerator control system based on the EPICS framework.

SYSTEM CONFIGURATION

The driver linac injector of the RAON consists of a 28-GHz superconducting Electron Cyclotron Radiation (ECR) ion source, the LEBT (low energy beam transport), the 500-keV/u RFQ (radio-frequency quadrupole) and the MEBT (medium energy beam transport). For the ECR ion source, superconducting magnets and dual high power RF sources of 28 GHz and 18 GHz are used to improve its performance [1]. The high voltage ion sources could get from two different high voltage platforms (50kV and 80kV). The Vacuum control system for the ECR ion source is consisted of Allen-Bradley PLC (AB PLC) modules. The AB PLC chassis consists of four chassis and are installed each of the electrical potentials racks. Each vacuum control devices are connected with AB PLC modules to control turbo pumps and to read pressure of the vacuum chamber. Vacuum gauge controller (XGS-600) is used to read pressure and to communicate with AB PLC through serial cable using RS232 protocol. Similarly, OSAKA turbo pump controller and LAYBOLD turbo pump controller are used to operate turbo pumps with AB PLC through serial cable. In order to construct the network system for connection among multi-voltage stages, we used remote IO modules 1756-AN2TR and 1734-AENTR of AB PLC. 175 6-AN2TR and 1734-AENTR modules are used to connect each of two racks on ground state through LAN cable by MOXA switch. And 1783-ETAP2F modules has used to connect among one rack of ground state and two racks installed on high voltage stages (50 kV and 80 kV) directly through optical fibers because communication failure occurred when LAN cable is used. The basic configuration of the control system is indicated by the network diagram shown in Fig. 1. Dashed lines are optical fibers and solid lines are LAN cables. Each chassis are installed at each platform as below Fig. 2. Internet Protocol (IP) address is assigned to two areas 192.168.1.* (area A) and 100.100.100.* (area B) according to voltage platforms level to reduce the risks from high voltage difference. The area 'A' is connected to total network of the test facility that included the ECR ion source facility. The area 'B' is local network that connects between remote IO modules of the AB PLC only. Because the IP address is not enough when configure the total network system of the test facility. The control system performs the interlocks for the vacuum system of the ECR ion source. And this system will be integrated with the Experimental Physics and Industrial Control System (EPICS) to operate the system record the parameter values by EPICS Input Output Controller (IOC) using "process variables" in real-time.

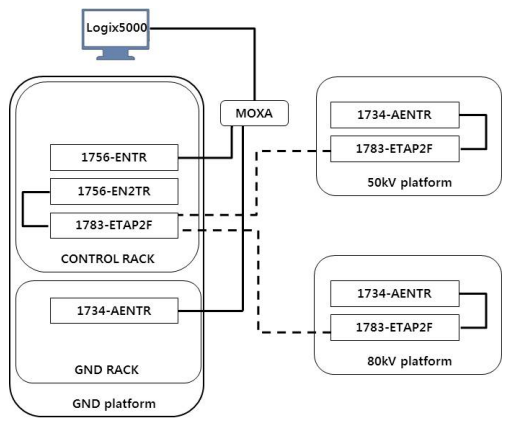


Figure 1: Diagram of network system for ECR ion source control system. 1783-ETAP2F modules for optical communication are used to connect between the high voltage platform and the control module that installed on ground platform.

VACUUM CONTROL SYSTEM

Vacuum control system of the ECR ion source consists of several turbo molecular pumps and controllers. They are installed at each voltage platforms. TC 353 turbo pump controller is installed at the ground platform and TC 2403 is installed at the 50kV platform. Both are made by OSAKA vacuum. TD 20 manufactured by LAYBOLD vacuum company is installed at 80kV platform. Each controller communicates with serial communication module (1734-232ASC) of AB PLC through serial cable. An AB PLC CPU 1756-L71 is used to control entire system of the ECR ion source and is installed at ground platform control rack of area 'A'

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

CAN OVER ETHERNET GATEWAYS - A CONVENIENT AND FLEXIBLE SOLUTION TO ACCESS LOW LEVEL CONTROL DEVICES

G.Thomas, D.Davids, CERN, Geneva, Switzerland
O.Holme, ETH Zurich, Switzerland

Abstract

CAN bus is a recommended fieldbus at CERN. It is widely used in the control systems of the experiments to control and monitor large amounts of equipment (IO devices, front-end electronics, power supplies). CAN nodes are distributed over buses that are interfaced to the computers via PCI or USB CAN interfaces. These interfaces limit the possible evolution of the Detector Control Systems (DCS). For instance, PCI cards are not compatible with all computer hardware and new requirements for virtualization and redundancy require dynamic reallocation of CAN bus interfaces to different computers. Additionally, these interfaces cannot be installed at a different location than the front-end computers. Ethernet based CAN interfaces resolve these issues, providing network access to the field buses. The Ethernet-CAN gateways from Analytica (GmbH) were evaluated to determine if they meet the hardware and software specifications of CERN. This paper presents the evaluation methodology and results as well as highlighting the benefits of using such gateways in experiment production environments. Preliminary experience with the Analytica interfaces in the DCS of the CMS experiment is presented.

INTRODUCTION

The control systems of CERN require the monitoring and control of a wide range of front-end equipment, much of which is installed in technically challenging locations with exposure to ionising radiation, strong magnetic fields and significant sources of electrical noise. A range of fieldbuses, including CAN bus [1], have been evaluated and selected by CERN to enable robust communication with the front-end hardware, leading to their widespread adoption in the LHC accelerator and experiments. As the readout computers and corresponding bus interfaces must be located outside of the hostile environment, the field buses are often more than 100 metres in length.

In the case of CAN bus, the communication on the bus is commonly driven by an OPC Data Access (OPC DA) server that runs on a computer and exposes the monitored hardware process variables and commands via the standard OPC interface protocol. The OPC server enables any OPC compliant client to connect and access the front-end data, independently of the details of the low level protocol used on the CAN bus for each type of hardware. A standard CERN solution for building the supervisory control system layer is to use the commercial SIMATIC WinCC Open Architecture (WinCC OA) [2] control system toolkit which natively supports OPC communication.

CURRENT STATUS AND LIMITATIONS

At the time of the construction of the LHC accelerator and experiment control systems, the connection between the control system computers and CAN buses was typically provided by internal PCI or PCIe cards or standalone USB connected interfaces. These CAN readout systems worked reliably throughout the first running period of LHC.

Changes in the computers running the control system software have been a significant motivating factor to re-evaluate the choice of CAN interface. For instance, the use of blade servers offers opportunities to minimize the rack space used by control system computers. The most compact blade servers do not provide the possibility to add PCIe expansion cards and therefore the PCIe-CAN interfaces are no longer a viable option in this context.

Due to the ubiquity of USB interfaces on modern computer hardware, the USB-CAN interfaces continue to provide a method of connecting CAN buses to any server. However, the limitation of maximum USB cabling length means that the computer and USB-CAN interface are typically installed either in the same rack or closely neighbouring racks. So with both PCIe and USB interfaces, the server and CAN bus interface must be co-located requiring that the control software servers are installed in or near to the racks where the CAN buses currently terminate. This imposes strong limitations on the possible re-organisation of rack hardware. With the current USB or PCIe interfaces, any significant change of server location would involve the costly installation of new CAN cabling.

Furthermore, there is increased interest in the use of redundancy at the control server level and the use of server virtualisation. Using the redundancy mechanism offered by WinCC OA, each control application runs on two computers. For robustness against power failures and network issues, it is beneficial to locate the two servers in different physical locations. As both computers need to access the front-end hardware, each server must be able to connect to the CAN bus interface, which makes it impossible to rely on the existing PCIe CAN or USB CAN devices used previously. Likewise, in the case of server virtualisation, the physical location of the server that runs the control software is not necessarily known in advance so it is essential to have a CAN interface which can be contacted from any location, independently of the distance between the server and interface hardware.

Offering network access to the CAN interface by Ethernet delivers a solution to these issues. This concept was demonstrated at CERN in the TOTEM and CMS experiments by coupling existing USB-CAN interfaces to Ethernet connected USB hubs. Using this approach the

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

STRIPPING FOIL DISPLACEMENT UNIT CONTROL FOR H- INJECTION IN PSB AT CERN

P. Van Trappen*, R. Noulibos, W.J.M. Weterings, CERN, Geneva, Switzerland

Abstract

For CERN's Linac4 (L4) Proton Synchrotron Booster (PSB) injection scheme, slices of the 160 MeV H- beam will be distributed to the 4 superposed synchrotron rings of the PSB. The beam will then be injected horizontally into the PSB by means of an H- charge-exchange injection system using a graphite stripping foil to strip the electrons from the H- ions. The foil and its positioning mechanism will be housed under vacuum inside a stripping foil unit, containing a set of six foils that can be mechanically rotated into the beam aperture. The band with mounted foils is controlled by a stepping motor while a resolver, micro-switches and a membrane potentiometer provide foil position feedback. The vicinity of the ionizing beam and vacuum requirements have constrained the selection of the above mentioned control system parts. The positioning and interlocking logic is implemented in an industrial Programmable Logic Controller (PLC). This paper describes the design of the stripping foil unit electronics and controls and presents the first results obtained from a test bench unit which will be installed in the Linac4 transfer line by the end of the 2015 for foil tests with beam.

INTRODUCTION

As part of the LHC Injectors Upgrade (LIU), CERN has planned the replacement of the current Linac2 by the Linac4 as injector to the PSB. This new linear accelerator (linac) is expected to increase the beam brightness of the PSB by a factor of 2, making a Large Hadron Collider (LHC) injectors' upgrade possible for higher intensity and eventually a luminosity increase [1]. A combination of bending, kicker and septum magnets will distribute the 160 MeV H- ions to the four superposed PSB synchrotron rings. The beam will subsequently be injected horizontally into the PSB using the stripping foil. The orbit of the circulating beam is displaced by ~81 mm, using two independent closed orbit bump systems, to meet the incoming beam [2]. The first injection bump uses four pulsed dipole magnets (BSW) magnets while a series of 4 horizontal kicker magnets (KSW), located in the PSB circulating orbit, will produce an additional closed orbit dump with varying amplitude to accomplish transverse phase space painting to the required emittance. This is illustrated in Fig. 1 [3]. The injected beam will be stripped by a H- charge exchange injection system (i.e. stripping foil) where the ions will be converted to protons. Partially stripped H- and ~1% H- missing the foil will be directed to an internal dump.

In order to reduce machine downtime it will be possible to remotely interchange the carbon foil (~1 μm thick) when deemed necessary, e.g. in case of rupture or decreased performance. For that reason a foil interchange mechanism (FIM) has been designed, incorporating a rotating band with six interchangeable foils. This paper discusses the electronic actuators, sensors and control system that make a foil exchange possible.

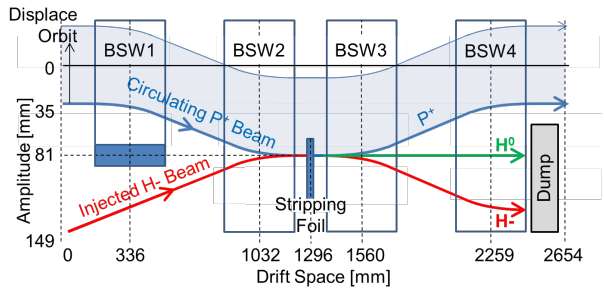


Figure 1: PSB injection region.

SENSORS & ACTUATORS

The FIM consists of a band, inside a vacuum chamber, that rotates over two pulleys so to displace a foil into the beam aperture. This will allow for a perpetual rotation so that each of the six foils can be reselected. Each foil is attached to a frame that has two important positions: *foil in* and *foil out*. Each of these positions allow for a 4 mm fine-tuning in order to find an optimum position for operation.

A retractable beam observation TV monitor (BTV) screen for beam position measurements on the selected foil is installed at the other side of the vacuum chamber, as shown in Fig. 2. When inserted it will be parallel next to *foil in* frame. To prevent BTV screen and FIM frame collisions, no FIM movement will be allowed when the BTV screen is in and at the same time any BTV movement is interlocked when the frame is not in the above mentioned positions. It is important to emphasise that a collision would result in major downtime because machine access and vacuum breakage will be required to replace the FIM. It is thus vital to know at all times the absolute position of the frame close to the beam aperture.

The choice of sensors is driven by obtaining that exact frame position and furthermore deal with the stringent vacuum pressure and radiation dose constraints. Furthermore the requested positioning reproducibility of 200 μm from the specification has to be taken into account.

Component Selection

The below electrical components are used to displace the band of ~50 cm so that one of the six foil frames is positioned into the injected beam aperture. The motor is the

* CERN TE-ABT-EC, pieter.van.trappen@cern.ch

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

A FAST INTERLOCK DETECTION SYSTEM FOR HIGH-POWER SWITCH PROTECTION

P. Van Trappen*, E. Carlier, S. Uyttenhove, CERN, Geneva, Switzerland

Abstract

Fast pulsed kicker magnet systems are powered by high-voltage and high-current pulse generators with adjustable pulse length and amplitude. To deliver this power, fast high-voltage switches such as thyratrons and GTOs are used to control the fast discharge of pre-stored energy. To protect the machine and the generator itself against internal failures of these switches several types of fast interlocks systems are used at TE-ABT (CERN Technology department, Accelerator Beam Transfer). To get rid of this heterogeneous situation, a modular digital Fast Interlock Detection System (FIDS) has been developed in order to replace the existing fast interlocks systems. In addition to the existing functionality, the FIDS system will offer new functionalities such as extended flexibility, improved modularity, increased surveillance and diagnostics, contemporary communication protocols and automated card parametrization. A Xilinx Zynq®-7000 SoC has been selected for implementation of the required functionalities so that the FPGA (Field Programmable Gate Array) can hold the fast detection and interlocking logic while the ARM® processors allow for a flexible integration in CERN's Front-End Software Architecture (FESA) [1] framework, advanced diagnostics and automated self-parametrization.

PRINCIPLE

In Fig. 1 a simplified schematic of a kicker system is shown. In general one can assume a system with matched impedances. A typical operation consists of charging the Pulse Forming Network or Line (PFN/PFL) to a voltage V by a Direct Current or Resonant Current Power Supply (DCPS/RCPS). When the main switch closes, a pulse of magnitude V/2 (if matched) propagates through the transmission line to the magnet. The pulse's energy is dissipated in the terminating resistor and full-field is established in the kicker magnet once the pulse reaches this resistor. Note that the magnet termination can also be a short-circuit, which doubles the magnet current because of reflection but also doubles the required fill-time of the magnet [2]. The pulse length in the magnet can be controlled by adjusting the timing of the dump switch relative to the main switch. When a sharp falling edge of the magnetic field is required for systems using a PFN, a clipper switch to ground is added between the PFN and the main switch. Typical values are 50 kV and 10 kA that traverse the magnet with pulse lengths between 500 ns and 600 µs.

During operation the system might be interlocked because of malfunctioning of one of the High Voltage (HV) switches. These malfunctioning causes are called *fast interlocks*, which

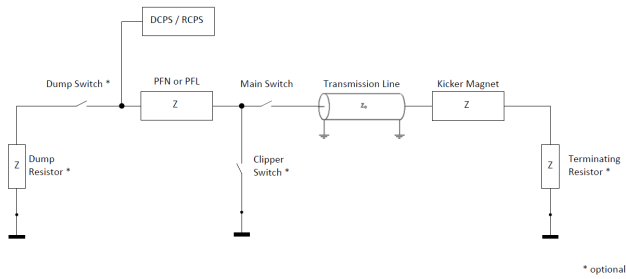


Figure 1: Simplified kicker system.

are detected over a large range of PFN voltages (i.e. the system's dynamic range). In addition to interlocking, they have often also a corrective action such as pulsing other magnets in order to fully deflect the beam and to avoid beam losses, i.e. to protect the machine. The following list gives a basic overview of the different types of fast interlock events:

- **Normal conduction** is not a fast interlock per se but for statistical reason it is useful to count the normal conduction cases.
- **Missing conduction** is registered when a trigger occurs but the switch is not pulsed. Lack of current in a HV switch or a missing load current after a certain amount of time after the trigger will result in a missing detection.
- **Erratic conduction** is the spontaneous conduction of an HV switch caused by tube or semiconductor malfunctioning as opposed to conduction following the normal triggering action. The FIDS mechanism registers an erratic when the switch conducts with no corresponding trigger pulse present.
- **Short-circuits** might happen at several locations, starting from the PFN up to a Terminating Resistor (TR). Depending on the magnet configuration several detection techniques exist, e.g. for short-circuited magnets one can compare the delay between the forward Main Switch (MS) pulse and the Dump Switch (DS) inverse current which, for normal operation, should represent the two-way transmission delay through the cable and magnet plus the single-way delay through the PFN.
- **No dump switch forward current** is the case when, for short-circuited magnets or systems with a Clipper Switch (CS) only, the short-circuit will reflect the travelling wave and the DS will have to withstand an inverse current (i.e. current from cathode to anode). In case of hollow-anode thyratron types, it is vital for the lifetime of the thyratron that it will first conduct forward current (i.e. anode to cathode) before the arrival of the inverse current.

* CERN TE-ABT-EC, pieter.van.trappen@cern.ch

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

UPGRADES OF TEMPERATURE MEASUREMENTS AND INTERLOCK SYSTEM FOR THE PRODUCTION TARGET AT J-PARC HADRON EXPERIMENTAL FACILITY

K. Agari#, Y. Sato, A. Toyoda, Y. Morino, KEK, Ibaraki, Japan

Abstract

Hadron experimental facility is designed to handle intense slow-extraction proton beam (750kW-15μA) from Main Ring of Japan Proton Accelerator Research Complex (J-PARC). On May 23, 2013, due to malfunction of the slow-extraction system in Main Ring, the production target in Hadron experimental facility was locally damaged because of overheat by absorbing proton beam in extremely short period (5 milliseconds). After the accident we have improved the monitoring system of temperature of the target with 100 milliseconds sampling rate in order to detect damage to the production target as soon as possible. The monitoring system has been operated without failure after the accident. This manuscript reports the upgrades of the temperature measurements at the J-PARC Hadron experimental facility.

INTRODUCTION

The Hadron experimental facility [1] (HD-hall) at Japan Particle Accelerator Complex (J-PARC) shown in Figure 1 is designed to handle intense slow-extraction proton beam from Main Ring (MR). The period of beam extraction from the MR to the HD-hall is 2 seconds and the operation cycle is 6 seconds.

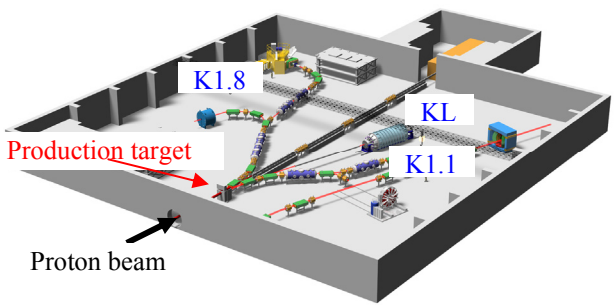


Figure 1: Illustration of the HD-hall.

On May 23th, 2013, 2×10^{13} proton beams were rapidly extracted to the HD-hall in 5 milliseconds due to malfunction of power supply of Extraction Quadrapole magnet in the MR. The production target [2] in the HD-hall was locally damaged because of rapid rise of temperature by beam deposit in extreme short period.

In order to detect the damage to the production target, the requirements of temperature measurement system are as follows.

- Upgrade read-out system of the production target temperature.
  - Hundred milliseconds sampling
  - Synchronization with beam extraction
- Make trend graphs and the waveform spectra of the target temperature as a function of time to tell the operators the state of the production target.

THE PRODUCTION TARGET

The production target at the HD-hall is made of gold and a copper block with coolant stainless pipes. Gold is chosen for high density, high thermal conductivity and good chemical stability. The total size of a gold structure is $15^{W} \times 6^{H} \times 66^{L}$ [mm]. The gold structure was divided into 6 pieces to reduce thermal stress. The water coolant pipes, embedded in the copper block, were made of stainless steel to avoid erosion and corrosion. The gold structure, the copper block and the coolant pipes are bonded by a Hot Isostatic Pressing (HIP) process.

The production target is located in an airtight chamber filled with helium gas. The beam entry and exit are covered with beam windows. Twenty thermocouples were attached at the surface of the gold pieces (12), the copper structure (2), the water and helium gas pipes (4) and the edge of the beam windows (2). The production target is designed to be capable for up to 50-kW proton beams.

MEASUREMENT AND CONTROL DEVICE

Before the accident we had measured the temperatures of the production target with 1-second cycle with a Programmable Logic Controller (PLC). The PLC had consisted of one EPICS-CPU, which is an embedded EPICS IOC on Yokogawa's FA-M3 PLC platform [3], three temperature monitor and one output modules. However, the temperature-measurement cycle of 1 second was too slow to detect the rapid temperature rise of the production target. Therefore we have upgraded the measurement system of the target temperature. Measurements of temperature with 100 milliseconds sampling are controlled with the sequence-CPU. The sampled data are stored in the sequent memory of the sequence-CPU. Then, an EPICS-IOC, running on the adjoining CPU module, installed next to the sequence-CPU, can take data from the sequence CPU via shared memory. Finally, the operators can monitor waveform spectra of the temperature rise on the production target, via the EPICS waveform record.

#agari@post.kek.jp

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

THE GENERAL INTERLOCK SYSTEM (GIS) FOR FAIR

F. Ameil, C. Betz, GSI Helmholtz Centre for Heavy Ion Research, Darmstadt, Germany
G. Cuk, I. Verstovsek, Cosylab d.d., Ljubljana, Slovenia

Abstract

The Interlock System for FAIR, called the General Interlock System (GIS), is part of the Machine Protection System which protects the accelerator from damage by misled beams.

The GIS collects various hardware interlock signals from up to 60 distributed remote I/O stations through PROFINET to a central PLC CPU. Thus a bit-field is built and sent to the interlock processor (PC) via a simple Ethernet point-to-point connection. Additional software Interlock sources can be picked up by the Interlock Processor via UDP/IP protocol.

The Interlock System for FAIR was divided into two development phases [1]:

- Phase A: interlock signal gathering (HW and SW) and a status viewer.
- Phase B: fully functional interlock logic with support for dynamic configuration, interface with Timing System, interlock signal acknowledging, interlock signal masking, archiving and logging.

The realization of the phase A will be presented here.

INTRODUCTION

The Interlock System for FAIR, called the General Interlock System (GIS), is part of the Machine Protection System which protects the accelerator from damage by misled beams.

The GIS will be used as a slower part of the Machine Protection System for all machines comprising the FAIR accelerator complex. Three components of the Interlock System can be identified [2] (Figure 1):

- the operation layer,
- the processing layer, and
- the signal pickup layer.

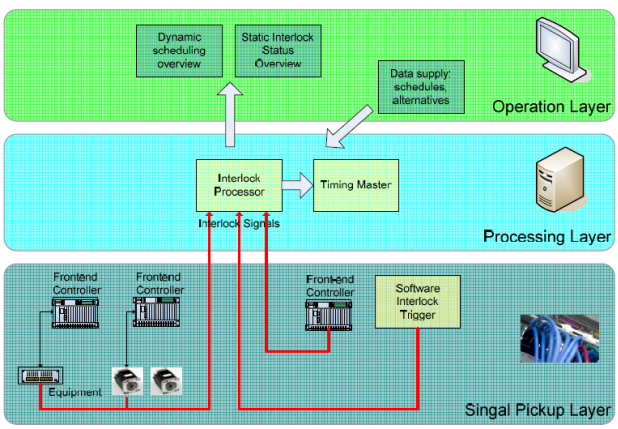


Figure 1: Layers of the interlock system.

The system must be able to collect hardware as well as software interlock signals from up-to 60 remote locations. Each remote station is capable of connecting to between 8 and 192 hardware interlock signals. Interlock transportation and processing must not take more than about 100 ms even when the system is extended to its limits - 4000 interlock signals. Based on these requirements, a combination of a PLC + interlock processor (PC) solution was chosen.

The purpose of Phase A was to confirm the selected architecture and components and to verify that such a solution is scalable and can reach the required performances. The solution will be installed and tested at the CRYRING accelerator in GSI.

The solution was developed by the company Cosylab d.d., Ljubljana, Slovenia, under supervision of the Control System Department at GSI.

SYSTEM DESCRIPTION

The main architecture of the GIS is presented in Figure 2 [3]. PLC Interface Modules (IM) with digital input modules are used to pick up hardware interlock signals. The control system network is used to pick up software interlock signals. After the GIS gathers both hardware and software interlock signals, it calculates summary interlock signals (specifics of this part will be defined and implemented in Phase B).

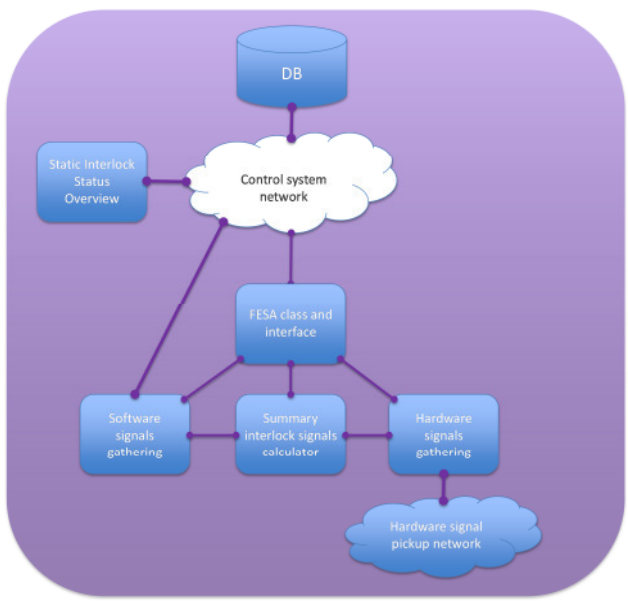


Figure 2: Architectural blocks of the GIS.

All changes of interlock signals are time-stamped. A user can use the GUI of the Static Interlock Status

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

A MODIFIED FUNCTIONAL SAFETY METHOD FOR PREDICTING FALSE BEAM TRIPS AND BLIND FAILURES IN THE DESIGN OF THE ESS BEAM INTERLOCK SYSTEM

R. Andersson*[1, 2], A. Monera Martinez[1], A. Nordt[1], E. Bargalló[1]
[1] European Spallation Source, Lund, Sweden
[2] University of Oslo, Oslo, Norway

Abstract

As accelerators are becoming increasingly powerful, the requirement of a reliable machine protection system is apparent to avoid beam-induced damage to the equipment. A missed detection of a hazard is undesirable as it could lead to equipment damage on very short time scales. In addition, the number of false beam trips, leading to unnecessary downtime, should be kept at a minimum to achieve user satisfaction. This paper describes a method for predicting and mitigating these faults, based on the architecture of the system. The method is greatly influenced by the IEC61508 standard for functional safety for the industry and implements a Failure Mode, Effects, and Diagnostics Analysis (FMEDA). It is suggested that this method is applied at an early stage in the design phase of a high-power accelerator, so that possible protection and mitigation can be suggested and implemented in the interlock system logic. The method described in this paper is currently applied at the European Spallation Source and the results follow from the analysis on the Beam Interlock System of this facility.

THE BEAM INTERLOCK SYSTEM

The Beam Interlock System (BIS) at the European Spallation Source (ESS) receives around 300 beam permit input signals and needs to stop beam operation in as short as 4-5 µs for some sections after detection of an error in the proton linear accelerator (linac) [1]. The requirement is based on possible damage from the high proton beam power on the surrounding sensitive equipment. The time scale requirement for the BIS is combined with a need to stop the beam in a reliable way, avoiding both false beam trips (stopping the beam without a hazard) and blind failures (missing to stop the beam when there is a hazard). By meeting these requirements, the BIS can help reaching the reliability goals for ESS and at the same time spare the equipment of any unnecessary damage.

There are four types of modules in the BIS at ESS. Together, they resemble a combined tree and star structure for a reliable transfer of the beam permit signal from the inputs to the actuators. Fig. 1 shows a conceptual flowchart of the BIS, disregarding the actual structure for simplicity. The BIS consists of close to 300 Fast Beam Interlock Driver (FBI_D) and Device Interface (FBI_DIF) Modules. For redundancy, each input signal goes to two FBI_DIFs. These modules interface with all input signals and transmit a beam permit signal to the BIS. Each signal into the BIS will be propagated through redundant links over the entire signal path. Table 1 displays nominal reaction times for the different modules in the BIS together with an example of the input signal from a Beam Current Monitor (BCM) and the LEBT chopper actuator. The table considers a worst-case scenario in terms of required reaction times. As is seen, if the modules perform as designed, the BIS achieves a total reaction time of less than 3 µs in the low energy part of the linac.

The rack configuration at ESS contains 24 enclosures with 36 racks each. All the signals from one rack enclosure are grouped into one pair of Master modules (FBI_M), making 48 FBI_M in total. The 48 FBI_Ms connect into two redundant Master of Master modules (MoM) located at the front end of the linac. These two MoMs then connect to three different Actuator Modules (FBI_A) – one for each actuator: Proton Source, LEBT chopper, and MEBT chopper. The FBI_As deliver the final trigger signal to the actuators to stop the beam, by deflecting beam to an absorber by the choppers, and inhibiting the creation of plasma in the proton source [2]. The BIS position in the Machine Protection (MP) conceptual level is seen in red in Fig. 2.



Figure 1: The BIS (red) transmits the input beam permit signals from the monitors and sensors to the actuators, which stop the beam in case of a hazard.

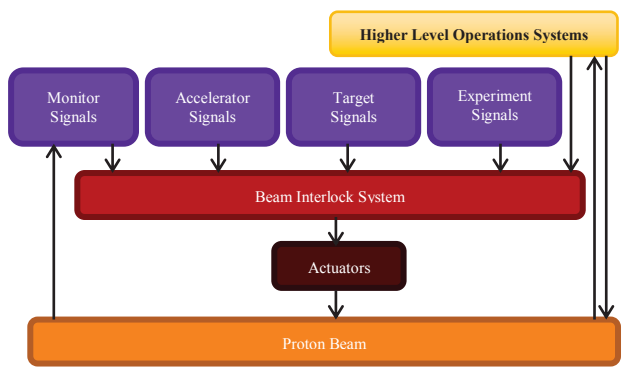


Figure 2: The BIS (red) role in the Machine Protection concept.

* riccard.andersson@esss.se

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

UNDERSTANDING THE FAILURE CHARACTERISTICS OF THE BEAM PERMIT SYSTEM OF RHIC AT BNL*

P. Chitnis#, T.G. Robertazzi, Stony Brook University, NY, USA
K. A. Brown, Brookhaven National Laboratory, NY, USA

Abstract

The RHIC Beam Permit System (BPS) monitors the anomalies occurring in the collider and restores the machine to a safe state upon fault detection. The reliability of the BPS thus directly impacts RHIC availability. An analytical multistate reliability model of the BPS has been developed to understand the failure development and propagation over store length variation. BPS has a modular structure. The individual modules have joint survival distributions defined by competing risks with exponential lifetimes. Modules differ in functionality and input response. The overall complex behavior of the system is analyzed by first principles for different failure/success states of the system. The model structure changes according to the type of scenario. The analytical model yields the marginal survival distribution for each scenario versus different store lengths. Analysis of structural importance and interdependencies of modules is also examined. A former Monte Carlo model [1] is used for the verification of the analytical model for a certain store length. This work is next step towards building knowledge base for eRHIC design by understanding finer failure characteristics of the BPS.

INTRODUCTION

The Beam Permit System (BPS) of Relativistic Heavy Ion Collider (RHIC) plays a key role in safeguarding against the developing faults in the collider. The energy stored in RHIC is about 2MJ in particle beams and 70 MJ in the form of magnet currents. Any abnormality in the machine can cause an undesirable escape of this energy, which can result into machine damage. BPS senses the onset of failures by monitoring the health inputs from RHIC support systems like power supplies, cryogenics, beam loss monitors, access controls, quench detection, vacuum system etc. In case of anomaly, it directs this energy out the machine for safe disposal.

RHIC reliability is directly dependent on the BPS reliability. Thus it is essential for the BPS to be highly reliable, being a safety critical system. Earlier a Monte Carlo (MC) model [1] was developed to quantify the probability of system level catastrophic events. This model simulated the behavior of modules with exponential competing failures. The module level failures were calculated by a quantitative fault tree analysis [2].

The MC model simulated the failure characteristics of the BPS for a certain store length. Also as the simulation takes about 17 hours to run 1E9 iterations to generate reproducible results. It generated probabilities of system level failures for a particular store length. This paper is aimed towards developing a model that generates the system failure probabilities as a function of store length. We develop an analytical model of the BPS behavior using stochastic mathematics for this purpose.

RHIC BEAM PERMIT SYSTEM

The BPS has a modular structure that consists of 33 Permit Modules (PM) and 4 Abort Kicker Modules (AKM). These modules are connected by three fiber-optic links that run 10MHz carrier links whose presence signifies that the system is healthy. These links are called the Permit Link (PL), Blue Link (BL) and Yellow Link (YL). The PMs collect support systems' health inputs from the field, called Permit Inputs (PI) and Quench Inputs (QI). The PM takes decision to drop the carrier, indicating a failure. When there is a magnet system failure, QI goes bad, and all the three links are dropped. In case of any other failure PI goes bad and only the permit link is dropped. The carrier drop propagates to all the modules, resulting respectively in magnet power +beam dump and a beam dump only. The AKMs send the beam dump signal to the beam abort system. They have redundancy incorporated. For more details of BPS refer to [3]. Figure 1 shows the BPS configuration.

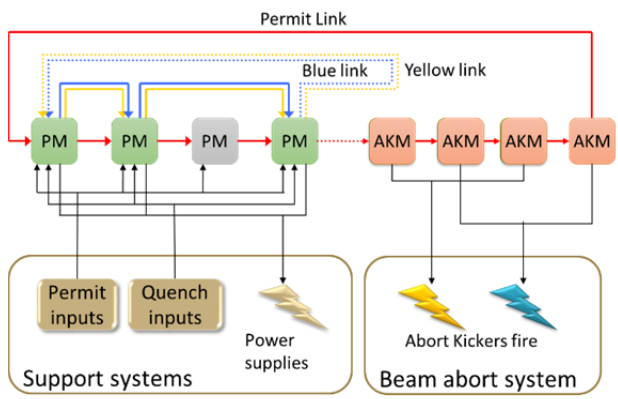


Figure 1: BPS connection diagram.

SURVIVAL DISTRIBUTION

The PMs and AKMs have been analysed for their failure modes in [2]. The top level failure of BPS depends upon the states of the PMs and AKMs. The calculation

*Work supported by Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy.
#prachi.chitnis@stonybrook.edu

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

INTERLOCK SYSTEM FOR MACHINE PROTECTION AT THOMX ACCELERATOR

N. Elkamchi, P. Gauron, H. Monard, LAL, Orsay, France

Abstract

ThomX is a Compton based photons source. It aims to produce a compact and directional X-rays source, with high performance, high brightness and adjustable energy [1]. The principal application fields are medical sciences, social technology and industry. An interlock system has been implemented for machine protection, especially to protect sensitive and essential equipment (magnets, vacuum system, etc.) during machine operation. It will allow taking appropriate actions to minimize recovery time in case of operation problem. ThomX interlock system is based on Programmable Logic Controller (PLC-Siemens S7-1500) distributed over the machine subsystems, it collects default signals from the different equipment of the machine, up to the central PLC which kills the beam, by stopping the RF or the injection, in case of problem (bad vacuum, magnets overheating, etc.). Actually, the interlock system is under implementation in parallel with the assembly of the machine. It will allow accelerator to work safely, and increase the availability of the machine.

INTRODUCTION

ThomX is a compact light source project where the Compton Effect is used to produce X-rays. A circular accelerator is combined to a high power laser, amplified in an optical cavity, to obtain collisions between an electron beam and the light pulses.

The collision between electrons and photons induces a diffusion of photons in the electrons direction with a high energy gain. Thus, from a few eV photons it is possible to obtain "hard» X-ray (a few tens of keV) using a 50 MeV accelerator. The machine presents the advantage of a compact dimensions (70 m2) and low costs.

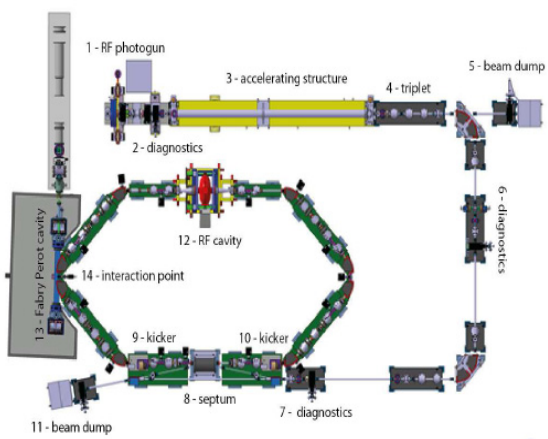


Figure 1: ThomX layout

Figure 1 shows the general layout of the machine. It is composed of four main systems: the injector, the storage ring, the laser and the Fabry-Pérot resonator.

The linac includes an electron RF photogun followed by an acceleration structure. The injection line consists of dipoles to steer the beam, quadruples and two diagnostic stations for energy spread measurements and beam characterization. The injection is performed by a septum and a fast kicker. [2].

The ring includes dipoles, quadruples, and sextuples for chromatic corrections. An RF cavity is used to bunch the electrons and to perform a longitudinal feedback.

The potential applications of this type of machine are numerous, including medical sciences, social technology and industry. Through collaboration between different French laboratories and industrials, it will be possible to share the technological excellence that will allow ThomX to be one of the most efficient machines in the world.

INTERLOCK SYSTEM DESCRIPTION

Because of the high beam power and energies in accelerator, the risk of beam induced damage is significant [3]. An interlock system is then used to protect the very sensitive and essential equipments (BPMs, vacuum chambers and valves, magnets) during machine operation. The ThomX interlock system is based on industrial programmable logic controller (PLC Siemens S7-1500 [4]). It collects default signals from the different equipment of the machine up to the central PLC which shut down the beam, by stopping the RF or the injection, when some parameters drifts outside specifications (bad vacuum, magnets overheating, etc.).

The interlock system consists of two levels. The first one is a local process, whose role is to monitor the variations of different parameters of the machine equipments, and generates a default signal in case of operation problem. The second level is the central PLC, which gathers and processes all the default signals from subsystems, and stops the RF power in a very short time.

To measure the performance of an interlock machine, we are based on two major parameters: reliability which measures the probability of fulfilling the major design function of the system, continuously and without interruptions, for a predefined period of time. The other parameter is the availability which defines the probability to find the machine fulfilling its major design function, when it is claimed to be in operation [5].

These two parameters describe the aptitude of the interlock system to meet the security requirements and to avoid failures.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

BUILDING AN INTERLOCK: COMPARISON OF TECHNOLOGIES FOR CONSTRUCTING SAFETY INTERLOCKS

T. Hakulinen, F. Havart, P. Ninin, F. Valentini
CERN, Geneva, Switzerland

Abstract

Interlocks are an important feature of both personnel and machine protection systems for mitigating risks inherent in operation of dangerous equipment. The purpose of an interlock is to secure specific equipment or entire systems under well defined conditions in order to prevent accidents from happening. Depending on specific requirements for the level of reliability, availability, speed, and cost of the interlock, various technologies are available. We discuss different approaches, in particular in the context of personnel safety systems, which have been built or tested at CERN during the last few years. Technologies discussed include examples of programmable devices, PLCs and FPGAs, as well as wired logic based on relays and special logic cards.

INTRODUCTION

There are three basic types of components in safety systems: *sensors* for collecting data on any measurable conditions important for safety, *actuators* for manipulating equipment important for safety when necessary, and *interlocks*, for computing the safety logic between the two. Sensors and actuators are always particular to the application in question, that is, the conditions to be surveyed and actions to be taken are different in each case. However, an interlock is just a unit for processing of logical information, and the technological choices for its implementation have usually more to do with requirements of overall throughput, reliability, cost, or technological diversity as mandated by principles of safety system design.

Safety interlocks are considered critical components that are subject to careful implementation and certification. International standards, such as IEC 61508 and 61511 for process industry, and IEC 61513 [1] for nuclear industry are often used in the design of entire safety systems, including the interlock. IEC 61508 and 61511 use the concept of safety integrity level (SIL) to quantify the required level of reliability of the safety system in its safety function as well as the ability of the system components to satisfy that requirement.

CERN GS/ASE group is responsible for all personnel safety systems, access control systems, and personnel safety alarm systems at CERN. We have a long experience of designing, building, and operating different personnel safety systems for CERN accelerators and experiments. Over the years, different approaches have been used for building interlocks for these systems, and this has given us some insight into the relative merits of the various technological choices. The technologies discussed in this paper are programmable logic controllers (PLCs), relay-based logic, wired logic with dedicated logic cards, and field-programmable gate arrays (FPGAs) with the examples based on actual implemented equipment.

PLC

Programmable logic controllers (PLCs) are the mainstay of modern process control applications. A PLC operates in a cyclic manner: during one cycle all inputs are collected, a new system state is computed in the CPU according to the program logic, and all outputs are set correspondingly. Cycle times of PLCs vary normally from a few to hundreds of milliseconds depending on the number of I/Os and the complexity of the program. I/Os are normally handled by separate modules, which can be placed some distance away from the CPU when connected via copper or fiber-optic cabling using a specific fieldbus protocol such as Profibus or Profinet [2]. Modern PLCs are also able to communicate via standard TCP/IP protocols via Ethernet, facilitating long-distance supervision of these systems. Special Supervisory Control and Data Acquisition (SCADA) systems are often used to integrate PLCs into a larger control framework.

Several manufacturers offer safety-related components certified for use in SIL-rated systems [3,4]. As a PLC-based system cannot be certified higher than SIL 3 (according to the standard, no system containing program code can be at SIL 4), that is normally the highest level of certification for the safety-related components as well. Some PLCs are able to integrate safety and non-safety-related programs and components within a single system, making communication between the two parts seamless.

At CERN, various generations of Siemens PLCs (S5, S7) are used in our personnel safety systems: LHC Access Safety System (LASS) [5], PS Access Safety System (PASS) [6], SPS Personnel Protection System, SPS Primary Ion Interlock [7]. This technology is fully mastered at CERN, and Siemens product life cycles are long allowing for a long utilization of installed hardware. Logic modifications and testing of PLC systems is fairly easy thanks largely to an integrated programming environment. Furthermore, PLC safety signatures provide a safeguard against unauthorized modifications of PLC code.

However, there are also drawbacks to PLC technology: For a large system, powerful CPUs are often needed to manage the logic coupled with sophisticated SCADA

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

DESIGN OF FAST MACHINE PROTECTION SYSTEM FOR THE C-ADS INJECTION I

Fang Liu, Jun Hu, XiaoShan Jiang, Qiang Ye, IHEP, Beijing, China
Guanghua Gong, Tsinghua University, Beijing, China

Abstract

In this paper a new fast machine protection system is proposed. This system is designed for the injection I of C-ADS which fault reaction time requires less than 20 us, and the one minute down time requires less than 7 times in a whole year. The system consist of one highly reliable control network based on a control board and some front IO sub-boards, and one nanosecond precision timing system using white rabbit protocol. The control board and front IO sub-board are redundant separately. The structure of the communication network is a combination structure of star and tree types which using the 2.5 GHz optical fiber links the all nodes. This paper pioneered the use of nanosecond timing system based on the white rabbit protocol to determine the time and sequence of each system failure. Another advantage of the design is that it uses standard FMC and an easy extension structure which made the design is easy to use in a large accelerator.

INTRODUCTION

The injection I of ADS (Accelerator Driven Sub-critical System) is being constructed in institute of high energy physics, China [1]. How to ensure the safe and stable operation of the injection, the fast protection system has played a crucial role. The protection system needs to make a quick decision on the current failure. When a general failure, the system needs to quickly adjust the system without stopping and set the working state to the normal mode. When a serious fault, it is necessary to quickly send out a signal and shutdown the whole system. The system also cans storage the statement before and after the fault. Through the storage of information the staff can quickly rule out the reason of occurrence failure [2].

FAST MACHINE PROTECTION SYSTEM FRAMEWORK

ADS fast protection system is composed of main control board, some front IO sub-boards, supervisory computer and one nanosecond precision timing system. The system framework is shown in Fig. 1.

Front IO sub-board monitors the 24V or 5V switch signal from the LLRF, PS, Vacuum, RFQ and other critical points. When the monitored points occur error, it sends the error signal to the main control board through optical fiber. The signal contains the time information that is labelled by the White Rabbit system. The main control board analyses and judgments the error through the stored fault tree. Then, it issued the processing results to the front IO sub-boards. Meanwhile, the main control board

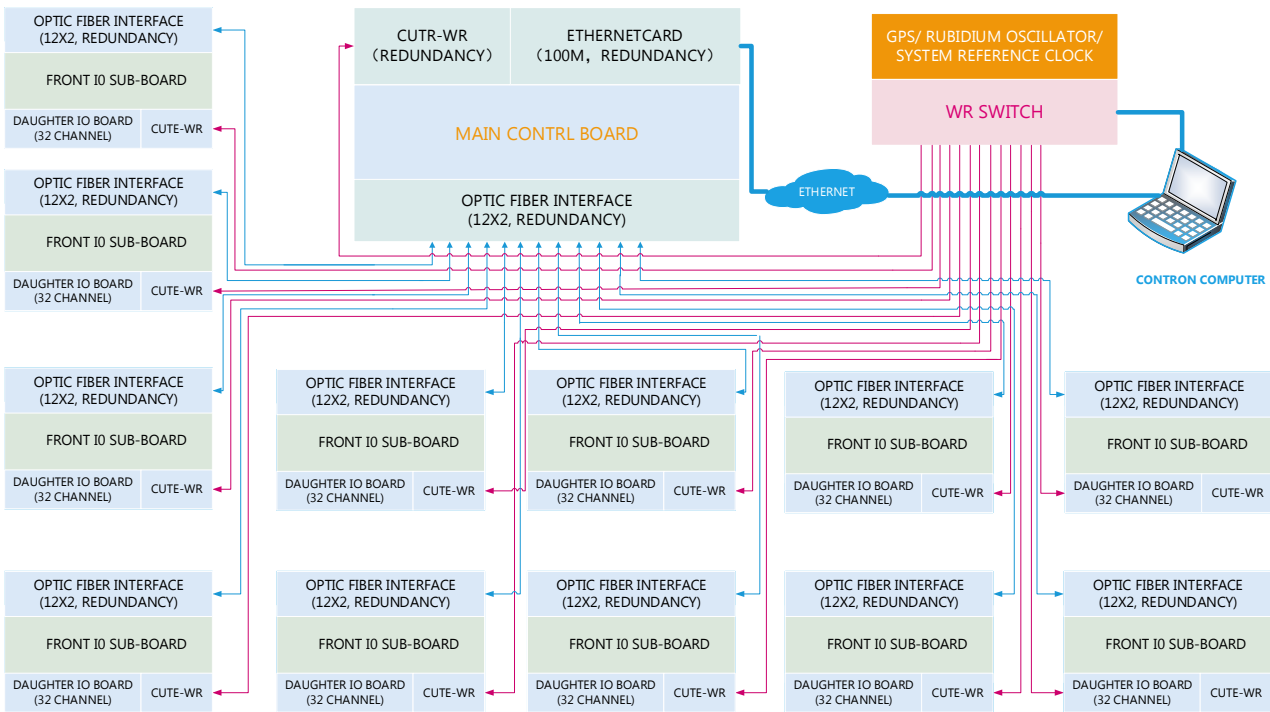


Figure 1: System Framework.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

UPGRADE OF THE TRIGGER SYNCHRONISATION AND DISTRIBUTION SYSTEM OF THE BEAM DUMPING SYSTEM OF THE LARGE HADRON COLLIDER

N. Magnin, A. Antoine, E. Carlier, V. Chareyre, S. Gabourin, A. Patsouli, N. Voumard,
CERN, Geneva, Switzerland

Abstract

Various upgrades were performed on the Large Hadron Collider (LHC) Beam Dumping System (LBDS) during Long Shutdown 1 (LS1) at CERN, in particular to the Trigger Synchronisation and Distribution System (TSDS): A redundant direct connection from the LHC Beam Interlock System to the re-trigger lines of the LBDS was implemented, a fully redundant powering architecture was set up, and new Trigger Synchronisation Unit cards were deployed over two separate crates instead of one. These hardware changes implied the adaptation of the State Control and Surveillance System and an improvement of the monitoring and diagnosis systems, like the various Internal Post Operation Check (IPOC) systems that ensure that, after every beam dump event, the LBDS worked as expected and is 'as good as new' for the next LHC beam. This paper summarises the changes performed on the TSDS during LS1, highlights the upgrade of the IPOC systems and presents the problems encountered during the commissioning of TSDS before the LHC Run II.

INTRODUCTION

LHC Beam Dumping System

The Large Hadron Collider (LHC) Beam Dumping System (LBDS) is responsible for the single turn extraction of the beam from LHC rings, after reception of a Dump Request (DR). It is composed, for each beam, of 15 extraction kicker magnets (MKD) followed by 15 extraction septum magnets (MSD) to extract the beam from the ring, followed by 10 dilution kicker magnets (MKB) in the extraction channel to spread the beam over the dump target (TDE) surface [1].

Each MKD and MKB is powered by a High-Voltage Pulsed Generator (HVPG), to produce the fast high current pulse needed to extract and dilute the beam, upon reception of a trigger.

Trigger Synchronisation and Distribution System

Within the LBDS control architecture, the Trigger Synchronisation and Distribution System (TSDS) is responsible for the detection of a DR from various sources, the subsequent generation of a dump trigger synchronised with the beam-free abort gap of the LHC, and the distribution of this Synchronous Beam Dump Trigger (SBDT) to the 25 HVPG of the 15 MKD and 10 MKB [2]. The TSDS is built around a major component, the Trigger Synchronisation Unit (TSU), which detects the DR from sources like the Beam Interlock System (BIS) [3], the State Control and Surveillance System (SCSS) [4], the Beam Loss Monitor Direct Dump (BLMDD), and issues a dump trigger synchronised with the Beam Revolution Frequency (BRF) signal [5].

The TSDS also includes an asynchronous Re-Triggering System (RTS) to cover the cases were the SBDT is not distributed properly, or when a MKD HVPG self-triggers [2].

During the LHC Run I operation, all beam dump requests were properly detected and distributed. Nevertheless some weaknesses of the TSDS were pointed out during various reviews and discussions, and a consolidation project was planned for Long Shutdown 1 (LS1).

Re-Trigger System Principle

The RTS is entirely passive and is composed of two redundant floating Re-Trigger Lines (RTL). A simplified view of one RTL is shown in Fig. 1. The 25 HVPG are interconnected by a chain of Re-Trigger Boxes (RTB), connected serially using a twisted pair cable. These RTB are composed of simple insulation transformers and decoupling diodes. When an HVPG triggers, five internal pickups in the HVPG will inject pulses in the RTL through five input transformers of the RTB. These pulses are then transmitted to all the other HVPG through two output transformers in each RTB. Thanks to this mechanism, in case of self-trigger of one MKD HVPG, all the other HVPG of LBDS will be triggered automatically as quickly as possible. Only the 15 MKD HVPG will inject pulses in the two RTL, as MKD see the circulating beam, and a self-trigger of one MKD HVPG requires retriggering all LBDS HVPG. The MKB HVPG do not inject pulses in the two RTL, because MKB do not impact the circulating beam, but a self-trigger of MKB HVPG will be detected by the various surveillance systems and a normal SBDT will be issued.

To cover a possible failure of the distribution of the SBDT issued by the TSU, an asynchronous dump trigger is also systematically generated by the TSU at the time it detects a DR. As up to 90 µs (one beam revolution) may be necessary for the TSU to issue a synchronous dump, the asynchronous dump trigger is delayed by 250 µs and injected in the RTL using a Trigger Delay Unit of 250 µs (TDU-250), as can be seen in Fig. 1.

The upgrades performed on the various TSDS subsystems, along with their motivations, are detailed in the following sections.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

ADaMS 3: AN ENHANCED ACCESS CONTROL SYSTEM FOR CERN

P. Martel, C. Delamare, G. Godineau, R. Nunes
CERN, Geneva, Switzerland

Abstract

ADaMS is CERN's Access Distribution and Management System. It evaluates access authorisations to more than 400 zones and for more than 35000 persons. Although accesses are granted based on a combination of training courses, administrative authorisations and the radio-protection situation of an individual, the policies and technicalities are constantly evolving along with the laboratory's activities; the current version of ADaMS is based on a 7 year old design, and is starting to show its limits. A version 3 of ADaMS will allow improved synchronization with CERN's scheduling and planning tools (used heavily during technical shutdowns, for instance), will allow CERN's training catalogue to change without impacting access management and will simplify and reduce the administrative workload of granting access. The new version will provide enhanced self-services to end users by focusing on access points (the physical barriers) instead of safety zones. ADaMS 3 will be able to cope better with changing and new requirements, as well as the multiplication of access points. The project requires the cooperation of a dozen services at CERN, and should take 18 months to develop.

INTRODUCTION

ADaMS is a system that centralises and standardises access policies throughout CERN; it is a central management system generating access authorisations for a multitude of different physical access systems. ADaMS has existed since 2007, and was designed for the 30 access zones existing at the time and that were exclusively related to safety, each with a relatively small area and with their access points geographically close to each other.

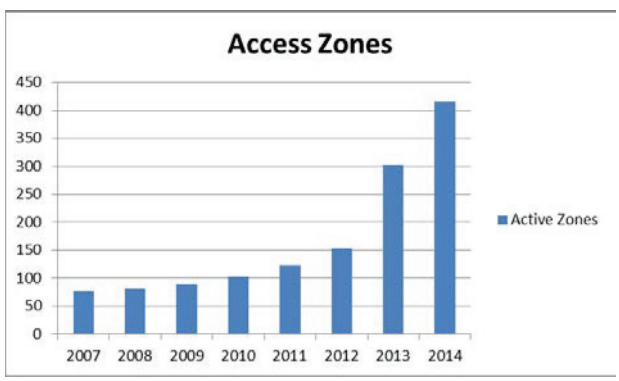


Figure 1: Evolution of the number of zones managed by ADaMS.

Today, the system is managing more than 400 zones (Fig. 1), not all related to safety, and with some of them concerning non-contiguous areas with access points very far from each other. At the same time, the business logic has changed dramatically both in quantity and quality; where originally only 5 criteria existed to grant or deny access (CERN access card, contractual situation of the person, dosimeter, followed courses and access request), new and more complex criteria have been added following new business requirements; IMPACT – a tool to coordinate the interventions on the accelerator premises during technical stops – now determines if access is granted or not, on top of the original criteria, for dozens of zones, and for around 200000 access authorizations (Fig. 2).

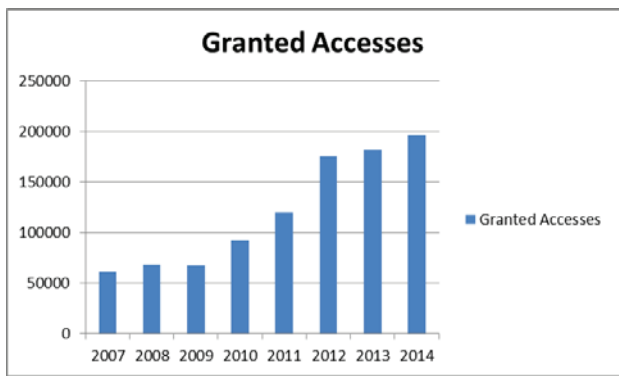


Figure 2: Evolution of the number of granted accesses by ADaMS per cycle.

With time, the changes and improvements to meet new requirements have been implemented on top of the original architecture; this means that even if for the past 8 years all new requirements were met with more or less complex adaptations of the original system, in some circumstances the system is no longer able to automatically justify the granting or revoking of an access right, as by design it does not collect and hold enough data for it. For these reasons, we have decided to build a new version from scratch.

MAJOR CHANGES FROM VERSION 2

Zones and Resources

One of the requirements that ADaMS 3 will have to fulfil is to be able to control access to entities other than "zones". Access to resources such as key cabinets, rooms

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

INTERLOCK OF BEAM LOSS AT LOW-ENERGY PART OF J-PARC LINAC

A. Miura#, Y. Kawane, N. Kikuzawa,
J-PARC Center, Japan Atomic Energy Agency, Tokai, Ibaraki, 319-1195, JAPAN
T. Maruta, T. Miyao
J-PARC Center, High Energy Accelerator Research Organization, Oho, Tsukuba, 305-0801, JAPAN

Abstract

The output beam power of the J-PARC linac has been improved by increasing the acceleration energy and peak beam current. The beam loss is getting serious along with increasing output beam power; however, the beam loss in the front-end region is difficult to detect because of the low energy of the radioactive emission. An interlock system using the beam current monitors has been developed to prevent significant material activation. In this system, an electrical circuit measures the beam transmission between the two beam current monitors. This study describes the design and performance of this electrical circuit and introduces the system configuration.

INTRODUCTION

The Japan Proton Accelerator Research Complex (J-PARC) linac can provide high intensity beams of peak current 50 mA, beam energy 400 MeV, pulse width 0.5 ms, and repetition rate 25 Hz using a radio frequency quadrupole linac (RFQ), three drift tube linac (DTL) cavities, 16 separation-type DTL (SDTL) cavities, 21 modules of annular coupled structure (ACS) cavities, and two beam transports, which have two ACS-type buncher cavities and debuncher cavities as shown in Fig. 1 [1, 2].

The J-PARC linac has a matching section between the RFQ and the DTLs, where 3-MeV beams are transported and measurements and beam profile matching are performed. An ion chamber-type beam loss detector has been employed for beam loss detection. The beam loss detector has a fast signal response; therefore, we have used this to develop an interlock system to stop beam operation when abnormal beam shots are detected. However, there is no beam loss detector in this section because the energy of radiation generated by 3-MeV beam is not sufficiently high to be detected.

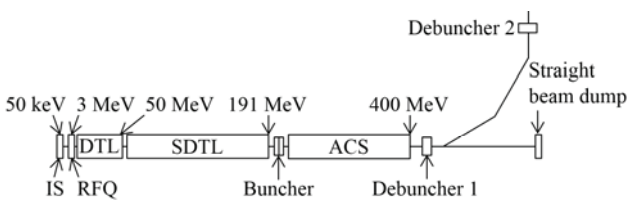


Figure 1: Schematic Layout of Present J-PARC Linac.

An interlock system that uses a new electrical circuit to measure the beam transmission between the two beam current monitors is proposed because of the large number of beam current monitors in the beam line.

In this study, a procedure for signal processing in the system and for processing waveforms in the processing unit is introduced. The waveforms obtained are also analyzed.

CIRCUIT DESIGN OF BEAM TRANSMISSION MEASUREMENT

The new interlock system, which uses beam current monitors, is named the beam transmission monitor (BTM) because the system compares beam currents detected by two independent beam current monitors. The signal design process is described below.

System Configuration

The schematic configuration of the system that uses the BTM is shown in Fig. 2. The system comprises two beam current monitors, and the signal processing circuit (BTM) shown in gray. Pre-amplifiers are usually used for continual beam operation and they are annually tuned to 0.10 V/mA of the beam current. The buffer amplifier has four channels per module, with two channels connected to the digitizer (one for both waveform outputs, 1 and 2) during operation, and the other two channels used by the BTM. The BTM circuit has six output terminals, as shown in Fig. 2.

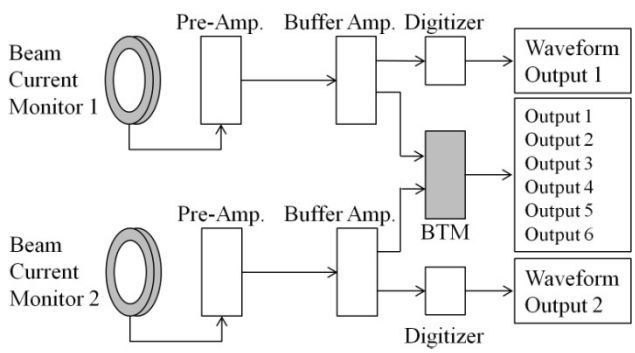


Figure 2: System Configuration of the Beam Transmission Measurement.

Process Flow

The signal processing flow in the BTM is shown in Fig. 3. It has two input terminals for the two signals coming from the two beam current monitors. Just after the signal input terminals, gain adjustment knobs amplify the input

#akihiko.miura@j-parc.jp

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

OVERVIEW AND DESIGN STATUS OF THE FAST BEAM INTERLOCK SYSTEM AT ESS

A. Monera Martinez, R. Andersson, A. Nordt, M. Zaera-Sanz, ESS, Lund, Sweden
C Hilbes, ZHAW, Winterthur, Switzerland.

Abstract

The ESS, consisting of a pulsed proton linear accelerator, a rotating spallation target designed for an average beam power of up to 5 MW, and a suite of neutron instruments, requires a large variety of instrumentation, both for controlling as well as protecting the different hardware systems and the beam. The ESS beam power is unprecedented and an uncontrolled release could lead to serious damage of equipment installed along the tunnel and target station within only a few microseconds. Major failures of certain equipment will result in long repair times, because it is delicate and difficult to access and sometimes located in high radiation areas. To optimize the operational efficiency of the facility, accidents should be avoided and interruptions should be rare and limited to a short time. Hence, a sophisticated machine protection system is required. In order to stop efficiently the proton beam production in case of failures, a Fast Beam Interlock (FBI) system with a targeted reaction time of less than 5 microseconds and very high dependability is being designed. The design approach for this FPGA-based interlock system will be presented as well as the status on prototyping.

INTRODUCTION

The team in charge of ESS Machine Protection is facing a great challenge: design, deploy and operate a Beam Interlock System (BIS) needed to protect the machine from beam-induced damage. The unprecedented beam power of 5MW at ESS can lead to severe damage within 10-20 microseconds only and thus beam operation has to be stopped reliably in an even shorter time upon the detection of non-nominal beam conditions. Such non-nominal beam conditions will be detected by several beam instrumentation systems such as the Beam Current Monitoring System, the Beam Loss Monitoring System and the Beam Position Monitoring System. The BIS shall fulfil demanding requirements regarding the hardware failure rate and beam availability as well as the requirements for the very short response time of a few microseconds. There are only four years to develop, test, produce and install the whole system. First commissioning of the ESS linac is planned for the beginning of 2018, and a first version of the BIS has to be ready by then. First protons on target are expected in mid 2019.

In a first step, the possible failure modes of the accelerator systems have been analysed. This analysis has been used to define around 170 different protection functions needed to protect the linac from beam induced damage, directly or indirectly. For each of these protection functions, an execution time has been defined as well as the required rate of dangerous failures. The highest rate for a function found is of 10^{-6} to 10^{-7} failures per hour with a response time of 5 microseconds (including the detection, processing and execution time) [1].

SYSTEM SPECIFICATIONS

In order to design a Beam Interlock System it is important to understand the system requirements and specifications. The design of the BIS is following the IEC61508 [2] and the IEC61511 [3] standards where applicable. However, since the BIS is a mission critical system and not a safety critical system, it is not needed to be compliant with those standards. Still, the team felt, that it is good to follow the guidelines provided by these standards in order to achieve a good level of protection thanks to the solid framework that the standards offer.

The following selected specifications allow to have a clear picture of what is expected:

Remote Monitoring: The BIS should be able to inform about its status and in case problems occur, operators shall be notified or an alarm shall be raised. For that it is required that the BIS will be connected to the main control system (EPICS) for monitoring and parameterization purposes.

Fast response time: the fast reaction time of the BIS is resulting from a time allocation which is done for a full protection function. A protection function includes time to detect a non-nominal condition of the machine, to process this information, to propagate and inform the BIS, process the information on BIS level and initiate a stop of beam operation. The fastest reaction time found for a protection function is 4-5 microseconds implying a reaction time of the BIS below 2-3 microseconds. This time differs for different beam energies along the linac and is only valid for the low energy part of the linac and can be relaxed for higher beam energies. Hence the Beam Interlock System will be called the Fast Beam Interlock (FBI) System. In order to achieve this reaction time the development of an FPGA-based system is required.

High availability: The FBI System should not interrupt the beam operation unnecessarily. One of the main requirements for the ESS facility is to deliver beam with very high availability. Thus, it is needed to keep downtime and false beam stops due to failures in the FBI System at a minimum. Hence the FBI System should be highly reliable as well as highly available.

Protection: The FBI System shall fulfil the requirement of a dangerous failure rate of 1.5 x 10^{-7} to 1.5

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

INTEGRATION OF PLC'S IN TANGO CONTROL SYSTEMS USING PyPLC

S.Rubio-Manrique, M.Broseta, G.Cuní, D.Fernández-Carreiras, A.Rubio,
J.Villanueva, ALBA-CELLS, Cerdanyola del Vallés, Barcelona, Spain

Abstract

The Equipment Protection Systems and Personnel Safety Systems of the ALBA Synchrotron are complex and highly distributed control systems based on PLC's from different vendors. EPS and PSS not only regulate the interlocks of the whole ALBA facility but provide a large network of analog and digital sensors that collect information from all subsystems; as well as its logical states. TANGO is the Control System framework used at ALBA, providing several tools and services (GUI's, Archiving, Alarms) in which EPS and PSS systems must be integrated. PyPLC, a dynamic Tango device, have been developed in python to provide a flexible interface and enable PLC developers to automatically update it. This paper describes how protection systems and the PLC code generation cycle have been fully integrated within TANGO Control System at ALBA.

INTRODUCTION

ALBA[1], member of the Tango Collaboration[2][3], is a third generation Synchrotron in Barcelona, Spain. It provides light since 2012 to users through its 7 beamlines,with 2 more under construction. The ALBA Control Section (ACS) is currently formed by 16 engineers devoted to the development of our Tango-based SCADA frameworks (Taurus[4], Sardana[5][6] and PANIC[7]) and PLC systems.

The ALBA Equipment and Personnel Protection Systems[8][9] (EPS and PSS) are distributed PLC-based systems autonomous from the Tango Control System. Both EPS and PSS are homogeneous systems based on single vendors (B&R and Pilz respectively). While the tasks to be done by the PSS are clearly specified and delimited by the ALBA Safety Group, the EPS PLC's became instead a versatile system that has been adapted for interlock, acquisition and motion control in both accelerators and beamlines.

Although other PLC based systems are used in ALBA to control the RF circulators, bakeout controllers and water or air cooling systems; the EPS is the most complex system managed by PLC's, using 58 B&R CPU's and 110 periphery cabinets to collect more than 7000 signals. In addition to the main purpose of protection, several hundreds of signals distributed across the whole system are acquired for diagnostics and control of pressures, temperatures and movable elements

The integration of the management of an independent system like EPS in the Tango Control System required of several phases, starting from the collection of cables from the Cabling Database to the final auto-generation of GUI's for both EPS Expert GUI and operator users (Taurus) .

GENERATION OF CODE VARIABLES

The Cabling and Controls Database

Every cable and equipment installed in the ALBA Synchrotron is registered in our Cabling and Controls Database (CCDB). Developed in 2007[10] by our Management and Information Software section (MIS) using MySQL and web technologies, it was the main support tool for the design and construction phase and now it is still kept updated as the main repository of equipments and configurations in our Accelerators and Beamlines. As of 2015 it lists 385 racks with 7131 equipments of different 874 equipment types. These equipments are connected using 20053 cables of 487 different cable types with a total length of 172.26 Km.

The CCDB provides, for each of our PLC CPU's or remote peripherals, the full list of connected devices, its equipment types, cable configurations, terminal used and the distribution of hardware in racks, including the routing of cables between hardware and control devices (Fig. 1). It also provides fast access to all available documentation for each type of equipment.

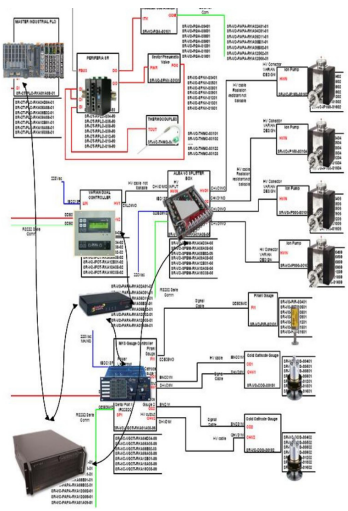


Figure 1: Diagrams of every subsystem have been produced prior to its introduction in our cabling database.

The PLC Auto-Generation Tool

The CCDB python API[11] provides full access to the Cabling Database from our control system tools. The API methods allow to search for equipments and get lists of connections, names and network information. These links are used to enable our Auto-Generation coding tool correlating the information of the equipments from the CCDB with the logics defined for them in the EPS.

To do so, all the common elements of our different Equipment Protection Systems have been standardized in

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

UPGRADE OF ABORT TRIGGER SYSTEM FOR Su erKEKB

S. Sasaki#, A. Akiyama, M. Iwasaki, T. Naito, T. T. Nakamura, KEK, Ibaraki, Japan

Abstract

The beam abort system was installed in KEKB to protect the accelerator equipment and the Belle detector, and for radiation safety, from high current beams. For SuperKEKB, the new abort trigger system has been developed. It collects more than 130 beam abort request signals and issues the beam abort trigger signal to the abort kickers. The request signals are partially aggregated in local control rooms located along the SuperKEKB ring and finally aggregated in central control room. In order to increase the system reliability, the VME-based module and the E/O module was developed, and all the abort signals between the modules are transmitted as optical signals. The system also has the timestamp function to keep track abort signal received time. The timestamps are expected to contribute to identify the cause of the beam abort. Based on the feasibility tests with a prototype module, the new module design has been improved and fixed. This paper describes the details of the new abort trigger system.

INTRODUCTION

SuperKEKB is the upgrade of the KEKB asymmetric energy electron-positron collider in Japan [1]. The designed luminosity is 40 times as high as the peak luminosity of KEKB.

SuperKEKB consist of two storage rings, low-energy ring (LER) and high-energy ring (HER). Each ring has a beam abort system to protect the accelerator equipment and the Belle detector, and for radiation safety, from high current beams. The abort trigger system collects more than 130 beam request signals and issues a beam abort trigger signal to abort kickers.

In KEKB, the modules of the abort trigger system were connected with twisted pair cable to transmit the electrical request signals and low pass filters were inserted to reduce electrical noise. The low pass filters caused time delay and the total system response time was about 100 μs. For SuperKEKB, the optical request signals are transmitted and the low pass filters have been removed. The response time is improved to be less than 20 μs.

The new system has the timestamp function to keep track of the abort signal received time. The resolution of the timestamp is 0.1 μs. It is expected to contribute to identify the cause of beam abort.

This paper describes the design and the status of the new abort trigger system.

#shinya.sasaki@kek.jp

SYSTEM DESIGN

Modules

The system is composed of two kinds of modules. First is 2ch E/O module which converts electrical signals from abort request source to optical signals. It can receive 3 kinds of electrical inputs: TTL, RS422, or relay, and can treat the input as active-high or active-low. They are set by slide switches in front panel. Figure 1 shows picture of the 2ch E/O module.



Figure 1: Picture of the 2ch E/O module. Input type and polarity can be set by slide switches (1424).

Second is VME-based 8ch beam abort optical input module (18K15). It gathers 8 optical signals and output an OR optical signal. Since it latches the input signals until reset, pulse input can also be detected and it keeps outputting once request signal is detected. Above functions are processed on FPGA, software is not running on the module. The module can be controlled and monitored via VMEbus. Figure 2 shows picture of the VME-based 8ch beam abort optical input module.



Figure 2: Picture of the VME-based 8ch beam abort optical input module (18K15).

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

DEVELOPMENT OF A NETWORK-BASED PERSONAL DOSIMETRY SYSTEM, KURAMA-micro

M. Tanigaki, Research Reactor Institute, Kyoto University, Kumatori, Osaka 590-0494, Japan
Y. Nakanishi, Electronics Laboratory, Shikoku Research Institute Inc.,
Takamatsu, Kagawa, 761-0113, Japan

Abstract

KURAMA-micro, a personal dosimetry system with network and positioning capabilities, is developed for the continuous monitoring of radiation exposure of individual in a large group, based on their action histories. Typical target users are the residents returning to their hometown after decontamination, and the workers involved in the decomissioning of the Fukushima Daiichi nuclear power plant. A KURAMA-micro unit consists of a semiconductor dosimeter and a compact DAQ board with a GPS module and a ZigBee module. Each unit records radiation data tagged with their measurement time and locations, and uploads the data to the server over a ZigBee-based network once each unit comes near the access point installed in the area of daily activities of users. Location data are basically obtained by a GPS unit, and an additional radio beacon scheme using ZigBee broadcast protocol is also used for the indoor positioning. A field test of the prototype of KURAMA-micro was performed in Fukushima city, and the radiation trends were successfully observed.

INTRODUCTION

The magnitude-9 earthquake in eastern Japan and the following massive tsunami caused a serious nuclear disaster of Fukushima Daiichi nuclear power plant. Serious contamination was caused by radioactive isotopes in Fukushima and surrounding prefectures. KURAMA [1] and KURAMA-II [2] were developed to overcome the difficulties in radiation surveys and to establish air dose-rate maps during and after the present incident. KURAMA has been successfully applied to various activities in the radiation measurements and the compilation of radiation maps in Fukushima and surrounding areas. KURAMA-II, an advanced version of KURAMA with autonomous operation and pulse height spectra measurement, has been used to establish the continuous monitoring scheme in residential areas in Fukushima prefecture. Fifty local buses and official cars equipped with KURAMA-II have been continuously operated in Fukushima prefecture, and radiation maps have been released to the public on weekly-basis. The results from KURAMA-II also contribute the predictions of air dose rates of the future in Fukushima.

Along with such wide-scale monitoring, the importance of more personalized monitoring intending to track the actual exposure of individual to the radiation are recognized to clarify the dependence of radiation exposures to human activities. For example, it is well known that the indoor radiation exposure becomes lower than outdoors, but the time length of stay indoors/outdoors vary depending on one's occupation, age, place to live etc. To clarify what kind of activities cause the higher radiation exposure is important to establish effective measures to minimize the radiation exposure. As the recovery from the nuclear accident in Fukushima progresses, strong demands of such personalized monitoring arise for a large group with higher risk of radiation exposure, such as the residents returning to their hometown after decontamination, or the workers involved in the decomissioning of the Fukushima Daiichi nuclear power plant.

For such purpose, personal dosimeters that measure the cumulative radiation dose, such as film badges, electronic pocket dosimeters, have been widely used, but the lack of positioning and time information made difficult to analyze in the view point of the human activities. Additionally, the difficulty in collecting data periodically from a large number of people obstruct the analysis of such data to extract the general tendency of radiation exposure towards human activities.

KURAMA-micro, a personal dosimetry system with network and positioning capability, is developed for such purpose. KURAMA-micro consists of a semiconductor dosimeter and a DAQ board with GPS and ZigBee modules. The data of KURAMA-micro is collected via an ad-hoc ZigBee based network to minimize the troubles of users. The development of a prototype KURAMA-micro is finished and field tests are on the way. In this paper, the outline of KURAMA-micro and the results from a field test in Fukushima are introduced.

SYSTEM OUTLINE OF KURAMA-micro

KURAMA-micro stands on a similar concept as claimed by KURAMA/KURAMA-II, and intends to be the alternative choice of conventional personal dosimeters. Therefore, the system must overcome the difficulties arising from its usage. For example, the unit must be comparable in size to other personal dosimeters, and the operating time should be long comparable to the replacement period of conventional film badges, typically one month. Therefore KURAMA-micro is designed to achieve better performance on power-consumption or sufficient compactness rather than on the detection efficiency or the measurement precision.

KURAMA-micro consists of the radiation detection part and the DAQ part (Fig. 1). The radiation detection part is a conventional accumulated radiation dosimeter with I2C connection to the DAQ part. The DAQ part has an I2C interface for the radiation detection part, a GPS unit for positioning, and a ZigBee unit for the data communication.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

INTEGRATION OF HETEROGENEOUS ACCESS CONTROL FUNCTIONALITIES USING THE NEW GENERATION OF NI cRIO 903X CONTROLLERS

F. Valentini, T. Hakulinen, L. Hammouti, P. Ninin
CERN, Geneva, Switzerland

Abstract

Engineering of Personnel Protection Systems (PPS) in large research facilities, such CERN, represents nowadays a major challenge in terms of requirements for safety and access control functionalities. PPS are usually conceived as two separate independent entities: a Safety System dealing with machine interlocks and subject to rigid safety standards (e.g. IEC-61508); and a conventional Access Control System made by integration of different COTS technologies. The latter provides a large palette of functionalities and tools intended either to assist users accessing the controlled areas, either to automate a certain number of control room operator's tasks. In this paper we analyse the benefits in terms of performance, cost and system maintainability of adopting the new generation of NI multipurpose cRIO 903x controllers. These new devices allows an optimal integration of a large set of access control functionalities, namely: automatic control of motorized devices, identification/count of users in zone, implementation of dedicated anti-intrusion algorithms, graphical display of relevant information for local users, and remote control/monitoring for control room operators.

INTRODUCTION

The GS/ASE group is responsible for the realization and the maintenance of Personnel Protection Systems ensuring the safety of personnel accessing to the various accelerator complexes of CERN. The group was responsible for the conception of a new generation of *access points* [3] to gradually replace the traditional rotating gates, some of them aged of few decades.

Even though the different system functionalities were clearly identified and specified since the beginning of the project, the realization strategy favored a construction by the integration of off-the-shelf (COTS) products rather than to privilege any form of custom development. This choice was mainly motivated by the limited internal resources, making it necessary to externalize the activity to a system integrator, and by the fact that the various functionalities to be provided were already present on the market with well-known and robust commercial solutions. Another motivation was the conviction that constructing a new system out of modular simple blocks would be more easy and rapid than to start a development from scratch. However, in our experience, it turned out that dealing with COTS can be quite a risky activity for a wide variety of reasons.

In [1], presented in the *International Conference on Software Engineering* of 1995, Garlan, et al., introduces the term of *architectural mismatch* to explain the different problems inherent in integration of COTS software components. Typically, many problems are caused by incompatibilities between the programming languages, the operating platforms or the database schemas of the products being integrated. The authors argue that at the origin of any architectural mismatch problem there is the fact that COTS components make several assumptions about the system architecture and its operational environment; consequently, when these assumptions conflict or do not match with each other, the simplicity of assembling complex systems from integration of COTS immediately disappears.

The biggest challenge for designing an access point turned out to be the integration of numerous heterogeneous components such as PLCs to command the movements of the motorized doors, *front-end* to deal with the real-time authorization management, *biometry* for authenticating users, *video system* for remote surveillance of the areas, *public address* system for broadcasting of audio messages, and many other devices. In this context our team is constantly seeking new solutions and technologies allowing to minimize the number of CPUs and information that need to be handled and exchanged between the different subsystems.

This paper will especially focus on the impact analysis that a completely COTS oriented realization strategy had on the architecture of the first series of our access points and it highlights the actual interest in the National Instruments technology [5] as a possible solution for the implementation of more rational system architectures.

THE ACCESS POINT CONCEPT

This section describes the ideas and the concepts behind any access point at CERN. Many would claim that an access point is simply made of "*one door, equipped with expensive position sensors, a video camera and a microphone allowing remote connections with guardians or operators*". However it is much more than that. An access point in a facility like CERN plays a crucial role in defending the installation against intrusions and accesses of people not holding the security requirements. At the same time, it ensures the safety of personnel every-day working inside the potentially dangerous areas.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

COMMISSIONING AND DESIGN OF THE MACHINE PROTECTION SYSTEM FOR FERMILAB'S FAST FACILITY*

A. Warner#, L. Carmichael, R. Neswold, J. Wu, N. Liu and D. Crawford FNAL Batavia, IL 60510, U.S.A

Abstract

The Fermilab Accelerator Science and Technology (FAST) Facility will provide an electron beam with up to 3000 bunches per macro-pulse, 5Hz repetition rate and 300 MeV beam energy. The completed machine will be capable of sustaining an average electron beam power of close to 15kW at the bunch charge of 3.2nC. A robust Machine Protection System (MPS) capable of interrupting the beam within a macro-pulse and that interfaces well with new and existing controls system infrastructure has been developed to mitigate and analyze faults related to this relatively high damage potential. This paper describes the component layers of the MPS system, including a FPGA-based Permit Generator and Laser Pulse Controller, the Beam Loss Monitoring system design as well as the controls and related work done to date.

INTRODUCTION

The FAST Facility comprises an electron injector based on the Advanced Superconducting Test Accelerator (ASTA)[1], a radio frequency quadrupole (RFQ) based proton injector and the Integrable Optics Test Accelerator (IOTA) storage ring. The electron beam is produced by a 1.3 GHz RF photo-injector and then accelerated to ~50 MeV by two 1.3 GHz SRF cryomodules, each containing a single 9-cell cavity. The beam will then be injected into the linear accelerator which consists of a 12-m long, 1.3 GHz 8-cavity superconducting cryomodule (CM2). This is a Tesla type III+ cryomodule[2] driven by a 5 MW klystron. The electron beam energy gain will be approximately 300 MeV at this stage.

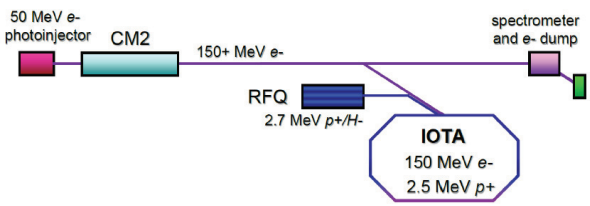


Figure 1: Schematic of FAST Facility.

The facility is also being expanded to accommodate further advanced accelerator research and development with the installation of a 2.5 MeV proton/H- RFQ accelerator. This accelerator starts with a 50 kV, 40 mA proton (or H-ion) source coupled to a pulsed 325 MHz RFQ to 2.5 MeV with a 1ms pulse duration for injecting into IOTA. This ring is 39 meters in circumference and will also be capable of storing electrons from 50 MeV to 150 MeV in energy. Figure 1 shows the placement of the ring in the FAST facility layout.

The Machine protection System (MPS) is being developed in stages that are commensurate with the commissioning goals for FAST. The primary objectives from the MPS point of view are to mitigate beam induced damage to the machine components and to provide a comprehensive over-view of the entire accelerator based on the input status of all the relevant subsystems [3].

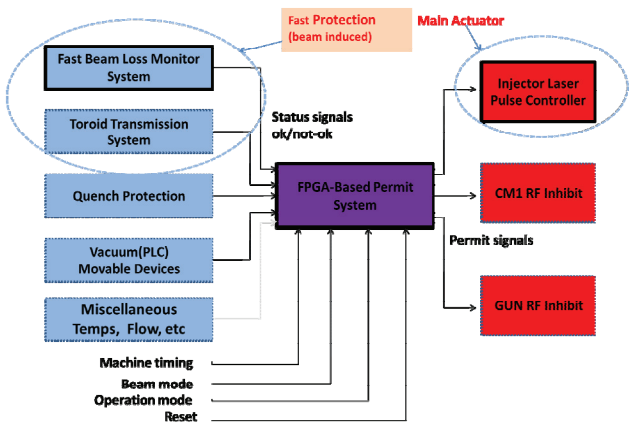


Figure 2: MPS Overview.

Figure 2 illustrates the overall MPS design which is divided into 3 layers; a sensor layer to collect sub-system status, a process layer that utilizes the status to generate the permits and an actuator layer to receive the permits and inhibit the beam. The initial stage of this development involved the design of the Laser Pulse Controller (LPC).

LASER PULSE CONTROLLER

The LPC is designed to be the primary actuator for beam inhibits. Its main function is to provide a gate to the gun laser system via the Pockels cells with a width that corresponds to the total number of 3 MHz pulses allowed without crossing the programmable threshold for losses. The maximum width of the gate is 1 ms which would accommodate a maximum of 3000 bunches. It is designed to inhibit the system within the 1ms macro-pulse window. The response time of the system has been demonstrated to be << 1 microsecond corresponding to less than 3 bunches in the machine. The LPC also provides a number of timing channels to the accelerator complex including a pulse which corresponds with the arrival of the first bunch; so-called first bunch trigger. This is an adjustable trigger delay with < 40 ps of jitter and tuneable to steps of

*Operated by Fermi Research Alliance, LLC, under Contract No. DE-AC0207CH11359 with the United States Department of Energy
#warner@fnal.gov

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

SAFETY INTERLOCK SYSTEM FOR A PROTON LINAC ACCELERATOR

Y. Zhao, Q.Ye, Y.Y.Du, J.He, F.Liu

Division of Accelerators, Institute of High Energy Physics, Beijing 100049, China

Abstract

The C-ADS Inject-I test facility is under construction in IHEP. An interlock system based on redundancy PLC is developed for machine protection and personnel safety. Device status, radiation dose, temperature of cavities and chambers are collected for machine state judge and interlock . A MPS(Machine Protection System) are working together with the interlock system in the control loop, and protect the machine in Four levels for different situations.

INTRODUCTION

The C-ADS linac include two Injectors and a main driver linac. Each injector is designed as a spare of another. The injector I proton linac in IHEP is consist of an ECR ion source, a LEBT (low energy beam transport line), a RFQ (radio frequency quadrupole accelerator), a MEBT(medium energy beam transport line) and superconducting spoke cavities[1]. The interlock system is designed and set up to protect personnel and devices from radiation hazards. It also will be a backup protection of MPS(fast Machine Protection System). The interlock system has permission signal, terminate beam signal and state signals collected from front device. Permission signals have the highest priority to access the Personal protection system, a key to the tunnel door will work together with the permission signals for safety. The state signals from Ion source, Power supply, LLRF (Low Level Radio Frequency) control system of RFQ, Vacuum and so on, indicate the corresponding system state or control “on and off”. The terminate beam signals will shut the Ion source down when emergency.

HARDWARE DESIGN

Phoenix RFC (Remote field controller) -460R has been chosen as the main controller for the interlock system. Two RFC can automatically build a high-performance connection to implement redundant through fiber optics for synchronization. Two Profinet interface in the RFC connected the controller into the Profinet ring with RSTP(Rapid Spanning Tree Protocol) to achieve fast ring detection, One LAN interface in the RFC is connected to the control network for. Before each cycle, One RFC which is set as the PRIMARY redundancy role will control the process，and transmits the data to another RFC which is set in Hot Standby mode as BACKUP redundancy role. If one Controller fails, the other will take over immediately [2].

The Axioline modules include bus coupler, I/O modules. The I/O modules complete the function signal in/output are connected to the bus couplers links to the Profinet ring, it’s the most fast I/O system in the world by now[2],the update time for each I/O modules is less than 1μs. The bus couplers has two Profinet interface which support PRL(The Phoenix Redundancy Layer), it will help to adopts valid values from I/O modules to the available Profinet controller.

The power supply of controller and Axioline modules are independent to ensure the reliability. Trio-diode modules are fixed for each two power supply units of the same type connected in parallel on the output side for redundancy. Those two power supply units will be isolated from one another. Figure 1 shows the redundancy diagram and Table 1 shows the signals from devices.

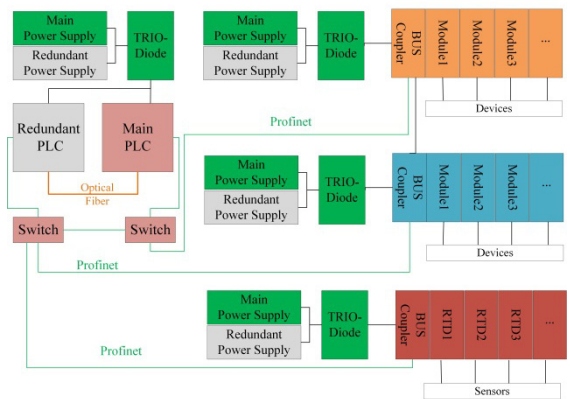


Figure 1: The redundancy diagram of interlock system.

Table 1: The Signals From Devices

Device	Input	Output
Ion Source	1	1
Vacuum	8	1
Personnel Protection	4	3
Power Supply	1	1
RFQ LLRF	1	1
Cryogenic	1	0
Super conductor LLRF	2	2
Bunching LLRF	1	1
MPS	1	6
Temperature	130	4
Flow meter	1	1

The fault tolerance requirements of C-ADS injector I are quite critical. So the reliability of the hardware is also considered. The index R (reliability), MTTF (mean time

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

REALIZATION OF A CONCEPT FOR SCHEDULING PARALLEL BEAMS IN THE SETTINGS MANAGEMENT SYSTEM FOR FAIR

H. Hüther, J. Fitzek, R. Müller, A. Schaller, GSI, Darmstadt, Germany

Abstract

Approaching the commissioning of CRYRING, the first accelerator to be operated using the new control system for FAIR (Facility for Antiproton and Ion Research), the new settings management system will also be deployed in a production environment for the first time.

A major development effort is ongoing to realize requirements necessary to support accelerator operations at FAIR. The focus is on the pattern concept which allows controlling the whole facility with its different parallel beams in an integrative way. Being able to utilize central parts of the new control system already at CRYRING, before the first FAIR accelerators are commissioned, facilitates an early proof of concept and testing possibilities.

Concurrently, refactorings and enhancements of the commonly used LSA (LHC Software Architecture) framework take place. At CERN, the interface to devices has been redesigned to enhance maintainability and diagnostics capabilities. At GSI, support for polynomials as a native datatype has been implemented, which will be used to represent accelerator settings as well as calibration curves.

Besides functional improvements, quality assurance measures are being taken to increase code quality in prospect of productive use.

COMMISSIONING OF CRYRING AT GSI

At the time of writing, the CRYRING heavy-ion storage ring, a Swedish in-kind contribution to the FAIR project, has been set up at GSI, with only few additional components still to be installed. The machine features an electron cooler, an RFQ linear accelerator and two injectors for different types of ions. See Fig. 1 for an overview of the CRYRING ring section.

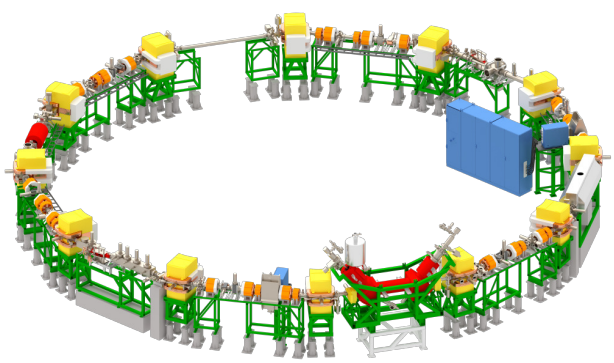


Figure 1: Overview of the CRYRING storage ring as setup at GSI (injection lines not shown), W. Geithner, GSI.

From a controls perspective, its commissioning represents a major milestone, as CRYRING will be the first machine solely operated via the new FAIR control system. Currently, the commissioning group still relies heavily on low-level control, but by the end of the year, they will have shifted their primary work place from the accelerator tunnel to the control room, utilizing the whole control system stack through the high-level applications provided.

TEST-BED FOR THE NEW FAIR CONTROL SYSTEM

Since CRYRING is equipped with its own injector line, it can continue operation even at times when the existing GSI accelerator chain is shut down for necessary FAIR upgrade and civil construction work, making it an ideal test-bed for the new control system. This way, it will contribute to validating concepts and technologies under real-world conditions [1], ensuring that the control system components work properly, individually and as a whole, and that business processes within and between involved parties are effective.

Although operating CRYRING does not imply the same requirements on the settings management system as the future FAIR facility will, core concepts necessary for highly flexible future operation scenarios can nevertheless be tested. As such, the beam-oriented approach to scheduling, designed for parallel beam operation at FAIR, will be utilized at CRYRING for the first time.

PARALLEL BEAM SCHEDULING CONCEPTS FOR FAIR

The designated operation modes of FAIR put demanding requirements on the new control system currently in development. To optimize the number of concurrent research programs, the facility will provide up to five beams in parallel with pulse-to-pulse switching between different particle types. Additionally, great flexibility shall be provided, allowing to change the parallel operation schemes on a daily basis.

Beam production chains and patterns are the central technical concepts within the new LSA-based settings management system to fulfill these requirements. Representing a major change in perspective, beam production chains establish a beam-oriented view on the facility, as compared to the accelerator-oriented view towards settings management dominant at GSI up to this point.

Beam production chains are defined using beam processes as atomic building blocks. Beam processes represent a specific procedure on the beam within one accelerator (e.g. injection, ramp, extraction). Within a beam production chain, the order of all beam processes necessary to provide the settings for producing a certain beam is described, from its

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

NUCLOTRON AND NICA CONTROL SYSTEM DEVELOPMENT STATUS

Evgeny V. Gorbachev, Vasiliy Andreev, Alexander Kirichenko, Dmitrii Vladimirovich Monakhov, Sergey Romanov, Tatyana Vladimirovna Rukoyatkina, Georgy Sergeevich Sedykh, Valery Volkov, JINR, Dubna, Russia

Abstract

The Nuclotron is a 6 GeV/n superconducting proton synchrotron operating at JINR, Dubna since 1993. It will be the core of the future accelerating complex NICA which is under construction now. NICA will provide collider experiments with heavy ions at nucleon-nucleon centre-of-mass energies of 4-11 GeV. The TANGO based control system of the accelerating complex is under development now. This paper describes its structure, main features and present status.

INTRODUCTION

NICA complex will consists of heavy-ion and polarized particles sources, RFQ injector, heavy- and light-ion linear accelerators, superconducting booster synchrotron, Nuclotron and two superconducting collider rings [1].

The control system of the NICA complex aims at accomplishing few main tasks:

- Management of large amount of equipment which is distributed on the accelerator complex area.
- Realization of different regimes of the accelerator complex working cycle – colliding or fixed target experiments, various ion types and energy.
- Strict synchronization of accelerators in the chain.
- Comprehensive beam diagnostics during the entire cycle.
- Providing protection and safety measures.

We can specify few important features that the NICA control system has to provide:

- Centralized administration and monitoring of the control system components.
- Reliable operation, quick recovery after possible failures.
- Access to equipment has to be restricted for certain personnel with rights limitation according to user roles.
- Ease of support, modification and scaling during long accelerator complex life-time which will operate until 2045.
- Rapid development and easy deployment of the control system, taking limited time and man-power into account.
- Possibility to integrate third-party control systems as some components of the accelerator complex will be designed and constructed by external organizations.

CONTROL SYSTEM LAYOUT

The control system is distributed network of computers which communicate by means of transport protocol over TCP/IP. The NICA control system uses the TANGO controls [2] as the middleware providing such communication protocol.

TANGO is the modern distributed control system framework based on CORBA. The fundamental unit of TANGO is a device, which is an abstraction hiding real equipment or program component behind the standard interface. TANGO provides high level client application interface which has necessary programming classes to implement client-server communications - synchronously or asynchronously execute commands, read or write attributes, or use events to acquire the data from the TANGO devices. TANGO incorporates a number of tools to build efficient control system environment including centralized administration and monitoring, access control, logging system, data archiving and code generation for rapid development of the TANGO device servers using C++, Java and Python.

Three layers of the NICA control system components can be distinguished (Figure 1):

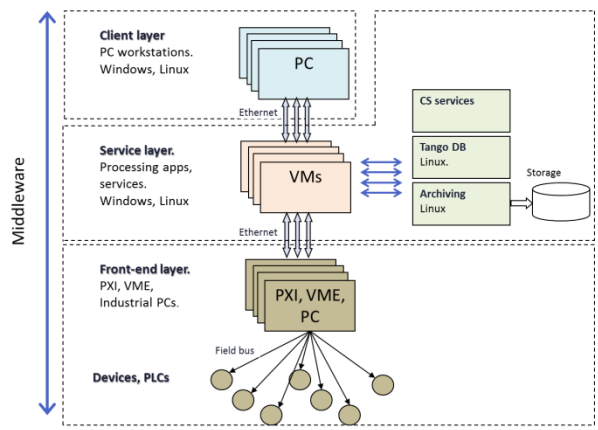


Figure 1: NICA control system layout.

1) Front-end layer consists of industrial computers, intellectual controllers, crates that directly control equipment and acquire data from sensors. Front-end computers run low-level TANGO programs that realize data acquisition, equipment handling and hide protocol and connection details from higher layer components.
2) Service layer consists of high level TANGO devices representing entire subsystems. They collect data from front-end TANGO devices, process it and realize algorithms to control some large subsystems. Those high-level programs provide standard TANGO interface for entire subsystems allowing client software to execute commands, read and write attributes without knowledge of subsystems structure. Besides, the service layer provides a set of services that are necessary for efficient control

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

IMPROVING SOLEIL COMPUTING OPERATION WITH A SERVICE-ORIENTED APPROACH

A.Buteau, B. Gagey, G.Abeille, Synchrotron SOLEIL, Gif-sur-Yvette, France
JC Fouquet, JCF company, Paris, France

Abstract

SOLEIL Computing division is in charge of managing computing infrastructures and applications on a 24/7 basis for SOLEIL staff and beamlines users.

During the last years SOLEIL accelerators and beamlines have been facing an ever growing dependency on Information Technologies to be able to deliver their service to their respective users. During the same period the Computing division had to manage an increasing number of technologies and software applications while continuously improving IT operation performances.

All this happened in a "classical" scientific environment where having immediate scientific results seems always more urgent than working on enhancing operational activities and minimizing the workload of IT groups.

This context has been the key driver to start the project of optimising our IT operational practices by adopting the ITIL [1] methodology with the following objectives in mind:

- Enhance the quality of IT services delivered
- Decrease the time spent by IT teams in operational activities to be able to keep resources focused on projects and development

The present paper will describe the overall vision of the project "Improving SOLEIL IT operation with a service oriented approach" and the strategy to make the methodology adopted efficiently by all IT groups.

THE MOTIVATIONS TO CHANGE OUR I.T OPERATIONAL PRACTICES

SOLEIL Computing Division Organisation

SOLEIL Computing division is composed of about 40 people organized in 4 groups:

- The ICA group is in charge of software development for Controls and Data Acquisition
- The ECA group is in charge of electronics for controls and data acquisition
- The ISG group is in charge of databases administration and enterprise applications (such as financial ones, technical documents management, etc.)
- The ISI group is in charge of IT infrastructures : networks, servers, file storage , desktops

Operational Practices of IT Groups are Very Different

Depending on their role in SOLEIL organization these 4 groups may have different internal customers:

- ICA and ECA groups are very close to "business" activities as they are working only for Accelerators and beamlines
- ISG is closer to SOLEIL Administration division
- ISI group activities are focused on providing IT infrastructures for all SOLEIL activities from offices to control systems

Moreover these groups are using different operational tools (see Fig. 1) to follow-up their daily activities:

- ICA uses JIRA [2] to manage software changes, users requests and development projects
- ECA, ISI use a CMMS [3] (*Computerized Maintenance Management System*) to track operational activities on the hardware components they are in charge of.
- ISG uses Redmine [4] to follow-up their internal software developments
- The "business" operational groups (*accelerators operators and Experiment Hall coordinators*) also use an ELOG system to track the I.T incidents which occur on the synchrotron facility

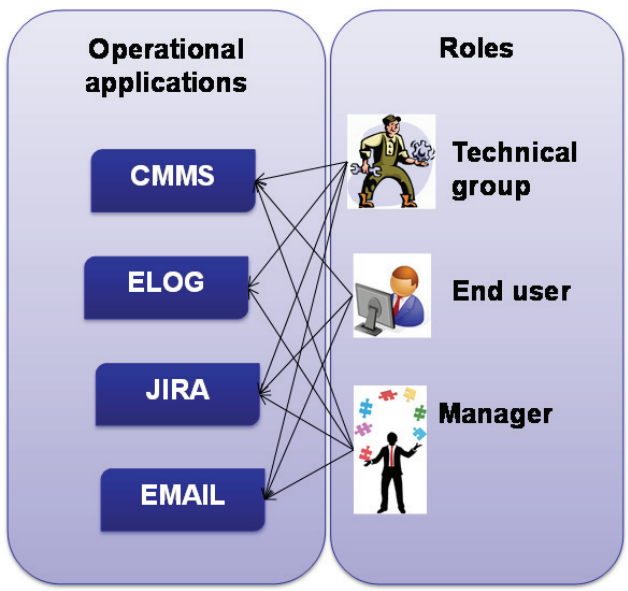


Figure 1: Operational tools used by various support groups.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

BEAM INSTRUMENTATION AND DATA ACQUISITION FOR CRYRING@ESR

T. Hoffmann, H. Bräuning, R. Haseitl, P. Miedzik, T. Milosic, R. Lonsing, A. Petit, A. Reiter
GSI, Darmstadt, Germany

Abstract

At FAIR the re-assembly of the well-known CRYRING accelerator, formerly hosted by Manne Siegbahn Laboratory (MSL) Stockholm, is currently in progress. This compact low energy heavy ion synchrotron and experimental storage ring will be a testing platform for all control system (CS) concepts decided on for FAIR. The CRYRING CS will be based on the system originally developed by CERN which combines the JAVA based application level LSA (LHC Software Architecture), the data acquisition level FESA (Front-End Software Architecture) and the White Rabbit based timing system. All parts have been enhanced with GSI specific functionality. In preparation for the commissioning of CRYRING later in 2015 all required beam instrumentation (BI) equipment including the software is now under development. We present the data acquisition (DAQ) concepts for the various instruments with emphasis on the seamless integration into the overall CS. For standard BI systems, such as digital video imaging, profile and intensity measurement, VME and IndustryPC based DAQ systems are used. For beam position monitoring a new hardware strategy which combines the microTCA and FMC (FPGA mezzanine card) form factors is under evaluation.

INTRODUCTION

The modernized CRYRING accelerator is presently being assembled in Experiment Cave B on the GSI Campus [1]. A schematic overview over site and accelerator design is given in Fig. 1. The accelerator consists of a 108 MHz 300 keV/u RFQ linac with a 50 kV MINIS ion source platform, the ESR injection transfer line and a synchrotron storage ring of 54.18 m circumference including an electron cooler and experiment section. Different singly charged ions in bunched or coasting beam will be provided by the linac. Later, the GSI experimental storage ring (ESR) will inject cooled, highly-charged heavy ions via fast extraction which can be accelerated up to 14.8 MeV/u in case of U^{92+} or 96 MeV/u protons. All accelerator parts are equipped with original beam instrumentation systems designed at MSL as well as new FAIR type solutions. The original MSL detectors and specialised pre-amps were preserved, but all DAQ systems have to be replaced.

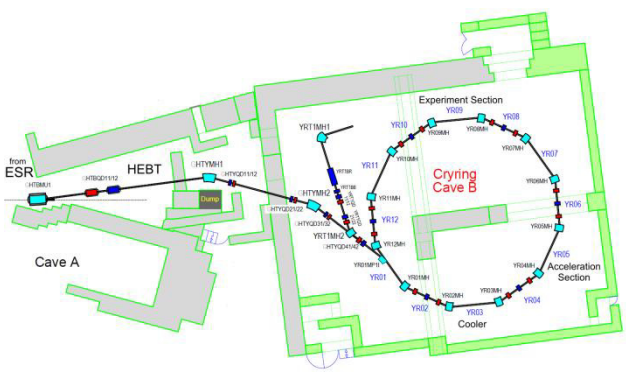


Figure 1: Scheme of CRYRING installation in Cave B.

The main intention in operating this accelerator, besides the physics aspects, is to provide a test platform with ion beams for many FAIR relevant systems in both hard- and software. Commissioning of the ion source has already started. The RFQ is ready for operation. The linac and the ring will be commissioned next year.

CONTROL SYSTEM

The control- and data acquisition system for CRYRING operation is based on a three tier architecture with FESA as the front-end level, the ZeroMQ based middleware and the top JAVA application level supported by JAPC (Java API for parameter control) and the LSA settings management [2]. CRYRING will be operated and synchronized with White Rabbit (WR) timing [3]. All BI DAQ systems need to be equipped with dedicated WR timing receivers in stand-alone, VME and PCIe form-factors. All those receivers were developed at GSI.

BEAM INSTRUMENTATION

CRYRING is equipped with a variety of beam instrumentation systems. For all of them new DAQ software has to be provided, which is also suitable for further use at FAIR. Presently the following BI systems are under development:

Intensity

A beam intensity measurement, mainly in the linac part, was realized by using Faraday-cups. The signals are digitized in a VME system by a SIS3302, 8-channel, 100 MSa/s, 16 bit ADC. The current amplifier copes with AC and DC beam and the DAQ calculates the total charge plus mean and maximum currents.

For AC beams automatic determination of the signal's baseline for every bunch is supported. For DC operation, the baseline may be measured on user request when no beam is hitting the cup. In both cases predefined values may be used in case of problems with the online measurement.

The region-of-interest to determine the mean current and total charge can be taken from two dedicated WR machine events indicating “beam on” and “beam off” or via user-defined markers. For offline analysis the raw

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

CURRENT STATUS AND PERSPECTIVES OF THE CRYOGENIC CONTROL SYSTEM OF EAST

Liangbing Hu, Ming Zhuang, Zhiwei Zhou, ASIPP, Hefei, China

Abstract

EAST (Experimental Advanced Superconducting Tokamak) is the first full superconducting experimental Tokamak fusion device in the world which has been carried out ten campaigns since its implementation at the end of 2005. The cryogenic control system for EAST was designed based on DeltaV DCS of Emerson Corporation which has been in operation for the same time period and has been proved to be safe and stable. However, Manny control components have been running beyond the expected lifetime gradually. Many problems from control system have affected the cryogenic system reliability. This paper presents the current status and upgrade solutions of the cryogenic control system of EAST.

INTRODUCTION

EAST (Experimental Advanced Superconducting Tokamak) is the first full superconducting experimental Tokamak fusion device in the world which has been carried out ten campaigns since its implementation at the end of 2005[1]. The cryogenic system is one important subsystem which is to cool down the superconducting magnets and relating components [2]. The total cold mass of EAST is about 250 tons. As shown in Figure.1, the cryogenic system is composed of a helium refrigerator and a helium distribution system. The helium refrigerator is composed of gas management system, compressors station, cold box and 10000 liter Dewar. All heat exchangers, absorbers and four turbine expanders are installed in cold box. The basic design capacity of the refrigerator is 1050 W/3.5K +200W/4.5K +13g/s LHe +12~30kW/80K [3]. To maintain the cryogenic state of the EAST cold components, the helium refrigeration system (HRS) supplies three helium coolants, SHe, liquid helium, and gaseous helium for the SC coils and their feeder lines, current leads, and thermal shield, respectively. The cryogenic distribution system (CDS) has sufficient mass flow to operate each of the SC coils.

Pulsed heat load is the main different factor between the cryogenic system of full superconducting Tokamak system and other large cryogenic systems. The cryogenic system operates in a pulsed heat loads mode requiring the helium refrigerator to remove periodically large heat loads in time. This operation mode must be taken into account for the design of the control system. The EAST cryogenic control system (ECCS) was designed based on DeltaV DCS of Emerson Corporation and has been proved to be safe and stable [4].

However, Manny control components have been running beyond the expected lifetime gradually. Many problems from control system have affected the cryogenic system reliability. This paper presents the current status and upgrade solutions of the cryogenic control system of EAST.

CONTROL SYSTEM OVERVIEW

As shown in figure2, the network of the cryogenic control system is composed of three parts: cryogenic redundant control local area network (LAN), data exchange LAN and main control data server LAN. There is a firewall between the intranet and extranet to ensure the security of the intranet. The cryogenic redundant control LAN employs a three-layer control structure including of the process layer, control layer and supervisory layer. The process layer includes all the field instruments, actuators, sensors and transducers. All the process variables and status information is converted to supervisory by control layer.

The control layer includes two local control cabinets for cold boxes and cryogenic distribution system and a remote control cabinet for compressor station. Local control cabinet includes redundant MD controller, power module and I/O cards. The remote control cabinet includes R5 remote I/O cards and serial card. The remote control system connects to local control by serial cards in terms of MODBUS protocol. New compressors are controlled by PLC and also connect to local control by MODBUS protocol. New PBS turbine is controlled by HEXTR provided by PBS Company and connects to DCS by local bus. New ATEKO turbine is controlled by PLC and as a profibus slave of DCS based on Profibus DP.

The supervisory layer is the interface between the operator and control system, including a professional plus station, four operator stations and an application station as OPC server. In this layer, engineers can configure the program software and manipulate for the cryogenic system.

The data exchange LAN is the auxiliary system of the cryogenic control system. The system includes database server, data acquisition and processing system, web server and FTP server.

The cryogenic data are stored in database server in data exchange LAN connecting to the cryogenic control LAN by application station using OPC protocol. Some important data is sent to the database of the EAST main control system. Some temperature of the magnet coils is acquired by technical diagnosis system. The temperature data are shown in the cryogenic monitor interface. So the process data will be transmitted to the database and

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

DESIGN AND STATUS FOR THE ELECTRON LENS PROJECT AT THE RELATIVISTIC HEAVY ION COLLIDER*

J.P. Jamilkowski#, Z. Altinbas, M.R. Costanzo, T. D'Ottavio, X. Gu, M. Harvey, P. Kankiya, R.J. Michnoff, T.A. Miller, S. Nemesure, T.C. Shrey, BNL, Upton, NY 11973, USA

Abstract

The Electron Lens upgrade project at the Relativistic Heavy Ion Collider (RHIC) has reached an operational status, whereby intense, pulsed or DC beams of electrons are generated in order to interact with the RHIC polarized proton beams in both the Blue and Yellow Rings at the 10 o'clock Interaction Region. Interactions between the electrons and protons are utilized to counteract the beam-beam effect that arises from the desired polarized proton collisions, which result in a higher RHIC luminosity. A complex system for operating the e-lens has been developed, including superconducting and non-superconducting magnet controls, instrumentation systems, a COTS-based Machine Protection System, custom Blue and Yellow e-lens timing systems for synchronizing the electron beam with the RHIC timing system, beam alignment software tools for maximizing electron-proton collisions, as well as complex user interfaces to support routine operation of the system. e-lens software and hardware design will be presented, as well as recent updates to the system that were required in order to meet changing system requirements in preparation for the first operational run of the system.

E-LENS GOALS AND DESIGN

The Blue and Yellow e-lenses are installed in the RHIC Ring at Interaction Region 10, in order to partially counteract the head-on beam-beam tune shift effect on the colliding RHIC beams, and thus permit RHIC proton beam operations at higher beam intensities, and therefore higher colliding beam luminosities for the RHIC experiments [1]. First commissioned during the FY2014 run, both electron lenses were successfully operated on a routine basis in a DC mode during the FY2015 RHIC 100 GeV polarized proton run [2].

The main components of each e-lens are the electron gun, electron collector, and superconducting solenoid magnet, though a set of additional systems is required for their routine operation. Beam transport magnets of both superconducting and non-superconducting varieties are controlled through two separate sets of standard VME hardware and software: warm magnets use equipment (PSC, QFG, and PSI) developed for the Injector machines within the Brookhaven National Laboratory (BNL) Collider-Accelerator (C-A) Department, and the cold magnets utilize the RHIC equipment (WFG, MADC) for reference control and readback information. Another system of note is beam instrumentation, which is primarily comprised of BPMs, current transformers, YAG crystal-based beam profile monitors, pinhole raster scan beam profilers, as well as new electron backscattering detectors [3].

TIMING SYSTEM

The Blue and Yellow e-lenses must be capable of operating in two distinct timing modes: pulsed mode for beam diagnostic measurements and system commissioning, as well as DC mode for routine operations. Activity in pulsed mode must be synchronized with the RHIC Event and Beam Sync Links, such that the timing system is capable of pulsing electron beam from the e-lens as a desired RHIC hadron bunch is passing through the Blue or Yellow Ring in IR10. To achieve these requirements, both e-lenses are dependent on the existing RHIC timing links as shown in Figure 1.

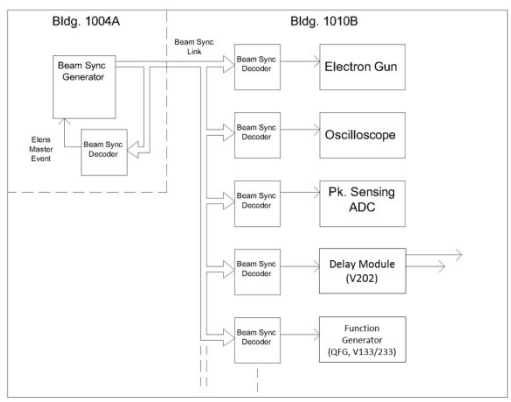


Figure 1: Schematic of the e-lens timing system, including RHIC components (building 1004A) and e-lens-specific components (1010B).

Software Interface

In order to enable the level of system flexibility required to operate the e-lens equipment under many different timing conditions while streamlining the user interface, a server-based ADO program is available to control the system that was written in the Java language. While it was not the first Java server written at the C-A Department, this software interface development project served as an important testbed for use of Java in developing ADO software interfaces rather than the standard C++ option.

*Work supported by Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy.
#jpj@bnl.gov

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

SIRIUS CONTROL SYSTEM: DESIGN, IMPLEMENTATION STRATEGY AND MEASURED PERFORMANCE

J.P.S. Martins, M. Bacchetti, E.P. Coelho, R.F. Curcio, J.G.R.S. Franco, R.P. Lisboa, P.H. Nallin, A.R.D. Rodrigues, L.D.S. Sachinelli, M.E. Silva – LNLS, Campinas, Brazil

Abstract

Sirius is a 3 GeV synchrotron light source that is being built by the Brazilian Synchrotron Light Laboratory (LNLS) [1]. The Control System will connect and operate all the equipment along the whole machine. The main goal of the Sirius Control System is to be distributed and digitally connected in order to avoid analog signal cables. A three-layer topology will be used [2]. The equipment layer uses RS485 serial networks, running from 6 to 12 Mbps, with a light proprietary protocol and CPU boards with proprietary hardware stacked on it, in order to achieve good performance. The middle-layer, interconnecting these serial networks, is based on Beaglebone Black single board computer and commercial switches. Operation layer will be composed of PC's running EPICS client programs. Special topology will be used for Orbit Feedback with a dedicated commercial 10Gbps switch. This paper will discuss the details of the Control System components, the implementation strategy for hardware and software and show some results of the prototypes.

INTRODUCTION

The Sirius Control System is designed to be scalable, distributed and easy to maintain. Currently, we are developing a generic solution for the hardware of the Control System – analog and digital interfaces, and also serial communications interfaces – to provide control features for the most of Sirius systems: vacuum system, pulsed power supplies, magnets power supplies, RF system, etc. The main characteristic of this solution is to be compatible with commercial low cost CPU boards. We achieve this using the SPI standard interface. The very first prototype for the Sirius Control System hardware was a homemade ARM CPU board with analog and digital modules, called PUC (Universal Control Board). Several tests have been done with this board and modules, and some results will be shown in this paper.

Besides the basic input/output, the hardware supports special features for synchronous operations, like curves for energy ramping of the booster and cycling magnets, and a circular buffer for post mortem acquisition.

With this strategy, the Controls Group will provide a low budget generic solution for hardware and software, which is reliable and of good performance in order to meet the requirements of operation of the various systems from Sirius.

HARDWARE IMPLEMENTATIONS

The main goal of the hardware solution provided by the Controls Group is to be easily adaptable to the equipment being controlled, avoiding the traditional "crate mounting" for a distributed installation, and sometimes inside the own equipment, preventing analog signals to travel over long distances.

Hardware Platform - GESPICON

This platform is under development and aims to be independent of a custom CPU board (like the PUC Base Board). Thus, any commercial CPU board, even the Beaglebone Black, can be used as embedded controller for the interface modules; it just has to have an SPI peripheral and some GPIOs. Figure 1 shows the components of the platform, called GESPICON (Generic SPI Controller). The connection of the stacked boards with the controller is done by a special bus, called SPIxxCON bus. This bus has the traditional SPI pins for serial data interchange, and some GPIOs for configuration of the transactions. It has also the pins for power supply and the optional use of quad-SPI features. The protocol of this bus is capable of addressing 8 module boards with 8 SPI devices in each one. Thus, one CPU controller can manage up to 64 SPI devices attached to it.

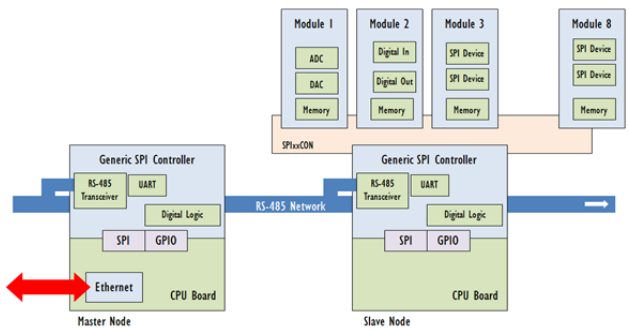


Figure 1: Schematic of the GESPICON hardware platform.

Another important feature of this implementation is the possibility to embed SPI flash memories on the interface modules. Thus, any interface can have static information used by the controller for auto-configuration. At boot time, the CPU module will search for each stacked module memories and download the data on it, which can be:

- Source code (Python or C) for the driver interface of the module;
- EPICS parameters for IOC configuration;
- Any parameter of the module devices;
- BSMP entities description (defined next page);

This system makes the CPU software independent of the interface modules drivers. Any module that would be developed in the future will be compatible with the CPU,

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

ARIEL CONTROL SYSTEM AT TRIUMF - STATUS UPDATE

R. Nussbaumer, D. Dale, K. Ezawa, K. Fong, H. Hui, M. Iranmanesh, J. Kavarskas, D. Morris, J.J. Pon, S. Rapaz, J.E. Richards, M. Rowe, T. Tateyama, E. Tikhomolov, G. Waters, P. Yogendran, TRIUMF, Vancouver, Canada

Abstract

The Advanced Rare Isotope & Electron Linac (ARIEL) facility at TRIUMF has now reached completion of the first phase of construction; the Electron Linac. A commissioning control system has been built and used to commission the electron gun and two stages of SRF acceleration. Numerous controls subsystems have been deployed including beamlines, vacuum systems, beamline diagnostics, machine protect system interfaces, LLRF, HPRF, and cryogenics. This paper describes some of the challenges and solutions that were encountered, and describes the scope of the project to date. An evaluation of some techniques that had been proposed and described at ICALEPCS 2013 are included.

SCOPE OF CONTROL SYSTEM

The e-Linac was to include a total of 3 stages consisting of 2 cavities per stage except for the first stage, which has only one, of Superconducting Radio Frequency (SRF) acceleration. The performance of the installed cavities has allowed postponing the third cryomodule. With the installation of the third cavity greater than 30MeV will be achievable. The second accelerator stage is planned for upgrade to 20MeV acceleration in one year from this writing. Deployment of the control system has been performed in subsystems, with the intention to isolate interaction between the subsystems, as well as to dedicate personnel to the various subsystems. Different broad control system technologies are used in each subsystem, and this allows personnel to focus on smaller classes of methodologies.

Numeric Facts and Figures

Devices under control	1160
Discrete IO Points	11155
EPICS IOCs	22
IOC Hosts	16
EPICS Device Support types	24

Figure 1: Scale of the control system at this writing

DEPLOYMENT STRATEGIES

A stated goal of the e-Linac control system was to re-use as much existing design and methodology as possible. The TRIUMF Controls Group has a significant body of work and methodology that was built up from prior projects. This includes a system of tools to build EPICS runtime databases, EPICS Operator Interface (OPI) screens, consistency checking in PLC interlock programming, and development of a consistent look and feel of control room consoles.

Part of the technology deployed in earlier TRIUMF control systems was based on consistent device styles chosen by other groups, especially vacuum, power supplies, and beam diagnostics groups. Despite the intent to develop no new methods and technologies, the style of control system interfaces used by equipment specialists from other disciplines did not accommodate the controls group strategy, and considerable re-work of existing methods was required. This paper will focus on some details of methods not previously included in control systems deployed at TRIUMF.

CONTROLS SUBSYSTEMS

Vacuum Systems

A new strategy for interface to vacuum systems was required when the vacuum group chose to introduce the concept of portable pumping carts. In this system, a small number portable carts containing turbo pumps, backing pumps, vacuum gauges and related valves was fabricated. These carts are arranged to pump down vacuum spaces at 26 discrete locations along the beamlines, according to stages of commissioning and following maintenance procedures requiring vacuum breakage.

Each turbo pumping cart is equipped with a small PLC drop that contains sufficient Input/Output support to drive and read the vacuum devices contained on the cart. A central vacuum controls PLC provides the program logic and EPICS interface for supervisory control and display. Asynchronous connection and disconnection of the central PLC is successfully accomplished by exploiting the ability of the Schneider[1] Quantum PLC to perform IO scans. IO scanning allows the PLC program to detect connect events and disconnect events seamlessly, and to map IO registers to beamline location-specific process variables.

The PLC reads dedicated loopback cables at each pumping cart and connection port, in order to permit the location of each connected cart to be tracked. This information is used to allow OPI displays to reflect the connection status of each cart and each beamline connection port in real time.

The ability to seamlessly connect and disconnect field IO at runtime has resulted in estimated saving of CDN $28000 in PLC hardware by removing the need for redundant dedicated PLC IO hardware. In addition, the reduction in cabling, estimated at 10000 metres, served to relieve the constraint of cable routing space availability to the Electron Hall where the accelerator and beamlines are

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

LANSCE CONTROL SYSTEM UPGRADE STATUS AND CHALLENGES*

Martin Pieck#, Dolores Baros, Eric Bjorklund, John Faucett, Jack Gioia, Jeff Hill, Pilar Marroquin, Jerome Paul, James Sedillo, Fred E. Shelley Jr., Heath Watkins
Los Alamos National Laboratory, Los Alamos, NM 87544, USA

Abstract

The Los Alamos Neutron Science Center (LANSCE) linear accelerator drives five user facilities: Isotope Production, Proton Radiography, Ultra-Cold Neutrons, Weapons Neutron Research, and Neutron Scattering. In 2011, we started an ambitious project to refurbish key elements of the LANSCE accelerator that have become obsolete or were near end-of-life. The control system went through an upgrade process that affected different areas of LANSCE. Many improvements have been made but funding challenges and LANSCE operational commitments have delayed project deliverables. In this paper, we will discuss our upgrade choices, what we have accomplished so far, what we have learned about upgrading the existing control system and what challenges we still face.

INTRODUCTION

The LANSCE accelerator looks back at almost 45 years of operations. It was one of the first accelerators to use computer technologies to control and monitor its beam line components. It started out with a custom in-house design called RICE (Remote Instrumentation and Control Equipment) which was installed in the early 1970's when the facility was built. Since then, the facility has seen partial upgrades and extensions utilizing CAMAC, VME, and PLCs while introducing EPICS (Experimental Physics and Industrial Control System) in the 1990's as a supervisory software control application.

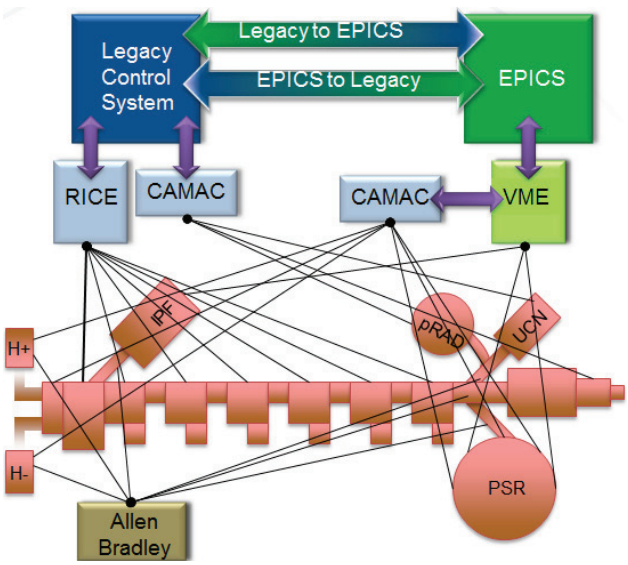


Figure 1: LANSCE Control System 1970-2010.

In 2011, LANSCE started a nine year rolling upgrade project which eventually will result in the complete replacement of the low-level RF system, the timing system, the timed data system, industrial control system, the beam-synchronized data acquisition system (including Beam Phase & Position Monitors (BPPM)), the fast protect reporting system and, wire scanner diagnostic equipment [1].

This upgrade project is primarily focused on ensuring reliable beam operations for a viable user program at the five experimental facilities. It will also benefit Los Alamos National Laboratory's future signature science facility called Matter-Radiation Interactions in Extremes (MaRIE) since, at its core, the 42-keV XFEL will be coupled with the existing LANSCE accelerator [2].

UPGRADE SCOPE

The monumental challenge of upgrading the control system is focused around the need to replace our VAX-based legacy control system which goes hand-in-hand with our RICE system. The VAX system has reached end of-life and the RICE systems, while a novel invention when it was designed in the late 1960's, is getting harder to maintain, and lacks the flexibility and performance of a modern distributed system with processing power near the front end.

RICE is a star configured control and data acquisition system that supports industrial control and beam-synchronized type of data acquisition. At its heart, the RICE Interface Unit can issue a parallel RICE module read request that provides a transverse snapshot of the accelerator. This implementation resembles the functionality of a timing system which provides trigger gates to distributed data acquisition equipment.

Given the complex functionality, replacing the RICE system is not an easy task and one starts to appreciate the engineer's ingenuity to design such a system about half a century ago.

Industrial Control

The first part of our RICE upgrade project addresses the Industrial Controls (slow control). We chose a Programmable Automated Controller (PAC) built by National Instruments (NI). The controller is called CompactRIO (cRIO) and is a reconfigurable control and data acquisition system which is supported by EPICS [3]. The NI cRIO–9024 embedded real-time controller features an industrial 800 MHz processor and contains, 512 MB DDR2 RAM, and 4GB of non-volatile storage. The removable cRIO controller sits in an NI cRIO-9118

*Work supported by LANL for the U.S. Department of Energy under contract W-7405-ENG-36; #pieck@lanl.gov ; LA-UR-15-27880

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

MaRIE – INSTRUMENTATION & CONTROL SYSTEM DESIGN STATUS AND OPTIONS*

Martin Pieck#, Robert W. Garnett, Brian G. Smith, Fred E. Shelley Jr.,
Los Alamos National Laboratory, Los Alamos, NM 87544, USA

Abstract

Los Alamos National Laboratory has defined a new signature science facility, Matter-Radiation Interactions in Extremes (MaRIE) that builds on the existing capabilities of the Los Alamos Neutron Science Center (LANSCE). It will be the first multi-probe materials research center to combine high-energy, high-repetition-rate, coherent x-rays with electron and proton-beam charged-particle imaging to perform in-situ measurements of a sample in extreme environments. At its core, a 42-keV XFEL will be coupled with the LANSCE MW proton accelerator. A pre-conceptual design for MaRIE has been established. Technical risk reduction for the project includes an injector test-stand that is currently being designed. New accelerators are either planned, under construction, or currently in operation around the world, providing opportunities for the MaRIE project to leverage the instrumentation & controls (I&C) efforts of these facilities to minimize non-recurring engineering costs. This paper discusses possible MaRIE I&C system implementation choices and trade-offs, and also provides an overview of the proposed MaRIE facilities and the current design.

BACKGROUND

X-ray imaging is unique, both because of the penetrating power of x-rays in solid matter - as Wilhelm Rontgen discovered in 1895 - and because x-ray wavelengths are short enough to resolve the interatomic spacing in matter via diffraction - Max von Laue's discovery in 1912. Those properties allow scientists to push forward fundamental physical sciences and to find major applications in structural imaging, from new commercial drugs to jet turbine blades [1].

Los Alamos National Laboratory's proposed Matter-Radiation Interactions in Extremes (MaRIE) experimental facility is slated to introduce the world's highest energy hard x-ray free electron laser (XFEL). The MaRIE 42-keV XFEL, with bursts of x-ray pulses at up to gigahertz repetition rates for studying fast dynamic processes, will help accelerate discovery and design of the advanced materials needed to meet 21st-century challenges [2].

MaRIE FACILITY

The MaRIE facility will include a 12-GeV linac to provide a suite of measurements designed to investigate the performance limits of materials in extreme environments. One of MaRIE's most powerful tools will be the ability to multiplex an x-ray FEL, electron, and proton radiography onto a target material to study dynamic events as they develop. The existing LANSCE proton linac will be used to provide proton radiography (pRad) [3]. The MaRIE electron linac will be built in a new tunnel north (right side in the figure) of the existing LANSCE proton linac tunnel as shown in Fig. 1 [4].

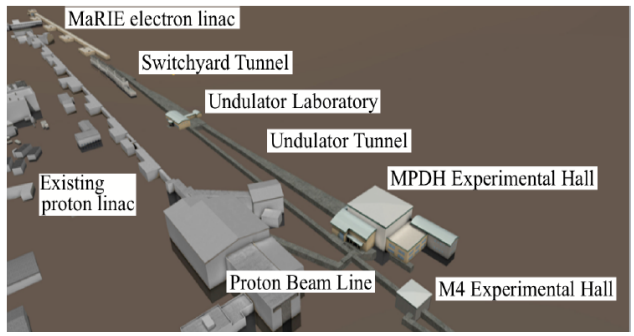


Figure 1: MaRIE Facility Layout.

OTHER XFELS

Besides the planned MaRIE XFEL facility in Los Alamos, New Mexico, USA, next-generation light sources also exist in Europe, Japan and elsewhere in the USA. X-ray facilities are being constructed at LCLS in California, SACLA in Japan, the European XFEL in Germany and the SwissFEL. The operating principles of these facilities are very similar. Electrons are first accelerated to high energies and then made to generate high-intensity x-ray laser light. LCLS and SACLA rely on conventional accelerator technologies. The European XFEL will use superconducting technology [5].

MaRIE BEAM REQUIREMENTS

The MaRIE electron beams consist of micro pulses for an XFEL undulator and micro pulses for electron radiography (eRad). A special feature of the MaRIE facility is the ability to provide unevenly spaced XFEL and eRad micro pulses distributed over a macro pulse of up to 100 µs. The macro pulse repetition rate is 60 Hz. Each XFEL micro pulse includes up to 0.2 nC of charge. Each 100-µs-long macro pulse can include up to 30 XFEL micro pulses. Each eRad micro pulse includes up to 2 nC of charge. Each 100 µs long macro pulse can include up to 10 eRad micro pulses. The spacing between micro pulses is determined by the experimental needs. The minimum spacing for each eRad micro pulse is 25 ns, while the minimum spacing for each XFEL micro pulse is 2.5 ns [4].

*Work supported by LANL for the U.S. Department of Energy under contract W-7405-ENG-36; #pieck@lanl.gov; LA-UR-15-27877

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

STATUS OF THE LOCAL MONITOR AND CONTROL SYSTEM OF SKA DISHES

S. Riggi*, U. Becciani, A. Costa, A. Ingallinera, F. Schillirò, C. Trigilio, INAF-OACT, Catania, Italy
G. Nicotra, C. Nocita, INAF-IRA, Bologna, Italy
V. Baldini, R. Cirami, A. Marassi, INAF-OATS, Trieste, Italy

Abstract

The Square Kilometer Array (SKA) project aims at building the world's largest radio observatory to observe the radio sky with unprecedented sensitivity and collecting area. In the SKA1 phase of the project, two dish arrays are to be built, one in South Africa (SKA1-Mid) and the other in Western Australia (SKA1-Survey). Each antenna will be provided with a local monitor and control system, enabling remote operations to engineers and to the Telescope Manager system. In this paper we present the current status of the software system being designed to monitor and control the dish subsystem. An overview of the dish instrumentation is reported, along with details concerning the software architecture, functional interfaces, prototyping and the evaluated technologies.

THE SKA DISH ARRAY

The Square Kilometer Array (SKA) project [1] is being designed to carry out radio observations of the sky in the wide frequency range from 50 MHz to 20 GHz with high sensitivity and collecting area to throw light in key areas of astrophysics and cosmology [2].

To this aim three different arrays were planned to be built for SKA Phase I and deployed at two different sites: the SKA1-Mid array at the South Africa's Karoo region hosting 190 dish antennas incorporating the Meerkat precursor telescopes and observing in the 0.35-13.8 GHz domain, the SKA1-Survey array at the Western Australia's Murchison region (MRO) site hosting 60 dishes equipped with phased array feeds (PAFs), incorporating the ASKAP precursor telescopes and observing between 0.65 and 1.67 GHz, and the SKA1-Low at the MRO site hosting 250000 low-frequency antennas.

The SKA design foreseen for Phase 1 was recently updated after a re-baselining process [3] carried out by the SKA Office to constrain the project costs to the established cost ceiling. The major outcomes of such review impacting dish array were a reduction in the SKA1-Mid array to 70% of its original size (190 dishes), equipped with prioritized band 2 (0.95-1.76 GHz), 5 (4.6-13.8 GHz) and 1 (0.35-1.05 GHz) and the deferral of SKA1-Survey array. Antennas (64) from the Meerkat precursor will be integrated in the SKA1-Mid array. The achieved array layout will have a central core component and additional spiral components targeted to deliver baseline lengths of 150 km. In the rest of the paper we will focus on SKA1-Mid antennas only.

* sriggi@oact.inaf.it

DISH INSTRUMENTATION

The general design for a SKA1-Mid antenna, shown schematically in Fig. 1, currently foresees a 15-m offset Gregorian dish with a feed-down configuration equipped with wide-band single pixel feeds (SPFs) for Band 1-2-5. The feed packages are mounted on a fan indexer at the focal position, allowing for changing among the available frequency bands. Only one band will be available for observation at any given time. A RFI-shielded cabinet is present in the antenna pedestal to house digital electronics and hardware for antenna movement and monitoring and control purposes, including the computing equipment supporting the Local Monitoring and Control system. No active cooling will be available for equipments inside the cabinet, only ventilation. Four sub-elements are identified in the SKA-Mid1 dish element: the Dish Structure (DS), the Single Pixel Feed (SPF), the Receiver (Rx) and the Local Monitoring and Control (LMC). Additional details on the first three are reported in the following subsections, while the LMC design is discussed in final sections.

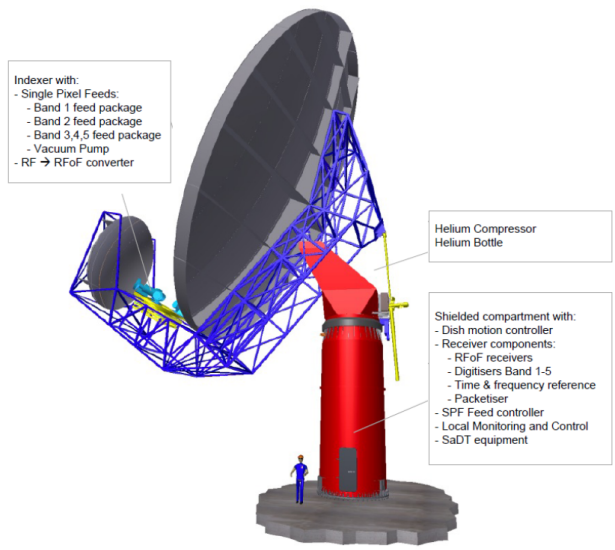


Figure 1: Schematic overview of SKA Dish design and instrumentation.

Dish Structure

The Dish Structure (DS) sub-element is responsible for the design of antenna, including mechanical structure and support, optics, reflectors, indexer as well as for the power distribution to be supplied to all sub-elements and safety systems (i.e. fire/intrusion sensors, limit switches, ...). DS

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

STATUS OF THE EPICS-BASED CONTROL AND INTERLOCK SYSTEM OF THE BELLE II PXD

M. Ritzert*, Heidelberg University, Germany†

Abstract

The Belle II e+e- collider experiment at KEK (Tsukuba, Japan) will include a new pixelated detector (PXD) based on DEPFET technology, providing the two innermost layers around the beampipe. This detector requires a complex control and readout infrastructure consisting of several on-sensor ASICs and remote FPGA boards. We present the architecture and EPICS-based implementation of the control, alarm, and interlock systems and their interconnectivity to other legacy/heterogeneous subsystems. The interface to the NSM2-based Belle II run-control to orchestrate the PXD startup sequence is also presented. An installation of CSS is used to implement the user interface. The alarm system uses CSS/BEAST, and is designed to robustly minimize spurious alarms. The interlock system consists of two main parts: a hardware-based system that triggers on adverse environmental (temperature, humidity, radiation) conditions, and a software-based system. Strict monitoring including the use of heartbeats ensures permanent protection and fast reaction times. Especially the power supply system is permanently monitored for malfunctions, and all user inputs are verified before they are sent to the hardware. The control system is embedded into a larger slow-control landscape that also incorporates archiving, logging, and reporting in a uniform workflow for the ease of daily operation.

CONTROL SYSTEM

The PXD control system is built on the control framework EPICS[1], which is widely used in high energy physics systems. The Eclipse-based suite Control System Studio (CSS)[2] is being used as the operator interface framework.

The EPICS system will be deployed in a dedicated physical network segment only accessible via a gateway using the Channel Access (CA) protocol [1] for load reduction on the EPICS Input/Output Controllers (IOCs)[3] and access logging. Access control lists (ACLs) enforced within the IOCs will be used to implement various access levels (read only, user, administrator).

In order to achieve a consistent view on operating conditions during the experimental run and for analysis after a failure, a major amount of identified key values have to be recorded and archived centrally in a database. This is to ensure that in normal as well as abnormal situations a consistent and synchronized view on the data and whole system is possible.

The system will be deployed on x86_64 servers running Scientific Linux as the operating system. To allow for fast, repeatable installation and controlled updates, the required software is packaged as RPMs.

The following list details major devices of the system and their means of control via EPICS:

- The data handling engine (DHE) [2, 3] receives the global trigger and timing signals and distributes them to the sensors. The internal FPGAs connect to the outside world via the IPbus protocol. IPbus is based on UDP/IP and can be implemented purely in HDL. An IOC implementing the IPbus protocol in the DHE-specific way has been written.
- The actual DEPFET detectors are controlled and read out via a number of ASICs that are accessible via JTAG from the DHE. The IOC for the DHE and the firmware running on the DHE include code to access the JTAG chain. Actual commands are again transmitted to the DHE via IPbus.
- The power supplies communicate with the control PC via the CHROMOSOME protocol, a fault-tolerant middleware [4] that also includes a heartbeat to quickly detect a lost connection on both sides. Another custom IOC has been written to interface these devices.
- The online selector nodes (ONSEN) receive tracking information from other subdetectors via the data concentrator (DATCON) and high level trigger (HLT) and uses it to filter out events not in proximity to a projected particle crossing in the PXD. The ONSEN devices are controlled with an IOC running on the embedded Power PC in the FPGA via a memory mapped interface.
- The ATCA and µTCA crates housing the ONSEN, DHH and DATCON systems are accessible via IPMI.
- Belle II run control communicates via the NSM2 protocol developed at KEK [5, 6]. A bidirectional gateway also developed at KEK is used to interface with the EPICS world.
- The environmental conditions inside the detector volume are monitored via Bragg grated fibres read out with an optical interrogator from Micron Optics. An IOC that reads data from the device at high frequency (several 100 Hz) is under development. The data will be made available to EPICS either as a single data points down-sampled from the raw device rate, or put through an FFT transformation to obtain a frequency spectrum, which is then available as a waveform record. The latter mode is especially useful to analyze the minute movements of the detector.
- The actual slow control of the IBBelle CO_2 cooling plant is implemented in a PLC located within the plant itself. The Modbus protocol is used to monitor its

* michael.ritzert@ziti.uni-heidelberg.de
† For the DEPFET collaboration.

[1] http://www.aps.anl.gov/epics/
[2] http://controlsystemstudio.org
[3] http://www.aps.anl.gov/epics/base/R3-14/12-docs/AppDevGuide.pdf

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

ACTIVE MAGNETIC BEARINGS SYSTEM UPGRADE FOR LHC CRYOGENIC COLD COMPRESSOR, RADIATIONS MITIGATION PROJECT (R2E)

P. Arpaia, University of Napoli Federico II, Naples, Italy; CERN, Geneva, Switzerland
M. Girone, University of Sannio, Benevento, Italy; CERN, Geneva, Switzerland
M. Hubatka, MECOS AG, Winterthur, Switzerland
M. Pezzetti, CERN, Geneva, Switzerland

Abstract

During the normal operations of the Large Hadron Collider (LHC), the high hadrons flux level induced several Single Event Errors (failure caused by a particle passing through a sensitive device) to the standard electronics installed underground. Such events perturbed LHC normal operation. As a consequence, a mitigation plan to minimise radiation-induced failures and optimise LHC operation was started: Radiation to Electronics (R2E) mitigation project. The full paper will deal with the mitigation problem for LHC/Point 8 equipment and the main improvements for the equipment in LHC/Point 4, with special focus on the controllers for the Active Magnetic Bearings (AMB) used in the IHI-LINDE cold compressors. A proven approach based on frequency response analysis to assess the cold compressor mechanical quality will be presented. The hardware and software design, implemented to increase the global reliability of the system, will be highlighted. A corresponding experiment protocol was developed at CERN in collaboration with the Swiss Company MECOS and the Italian Universities of Sannio and Napoli Federico II. Preliminary experimental results showing the performance of the proposed approach on a case study for the cold compressor stage 1 in Point 4 will be finally reported.

INTRODUCTION

The level of the flux of hadrons with energy in the multi MeV range expected from the collisions at the interaction Points 4 and 8, which are not in the center on the cavern, unlike Point 1,2 or 5, (Fig. 1) will induce Single Event Errors (SEEs) in the standard electronics present in much of the control equipment. Furthermore, a risk of SEEs induced by thermal neutrons cannot be excluded. In the long-term, such events could perturb the LHC operation, possibly leading to critical situations for the machine elements and subsequent important downtime [1].

In particular this paper is focused on the impact on the electronics of MECOS Active Magnetic Bearing controllers. The high-performance bearing are applied in the LHC cryogenic cold compressor equipment in order to satisfy the need of covering vibration, stability and robustness in nominal operation and transition phases. In fact, magnetic bearing-equipped machines are suitable for unlimited, reliable and safe operation even in the presence of "large" residual unbalance levels [2]. Those high performances requires deeper user understanding of the magnetic bearing technology and a commissioned system cannot be modify so easily. The R2E Mitigation Project foresees shielding or relocating the equipment sensitive to radiation, which is presently installed in these critical areas, into safer areas. The majority of these mitigation activities were performed, in three phases, between (date) in Point 4 and between (date) in point 8. This document reports on these mitigation activities and their associated improvements. It presents the strategy applied, the enhancements on the modification for the new version of AMB used in the IHI-LINDE cold compressors which were needed to be replaced within a longer distance (from 20 m to 45 m).

STRATEGY APPLIED

The first LHC run took place from 2010 to 2013, on 14th February 2013 the Long Shutdown 1 (LS1) started. During the LS1 preventive maintenance, upgrade and necessary repairs were performed.

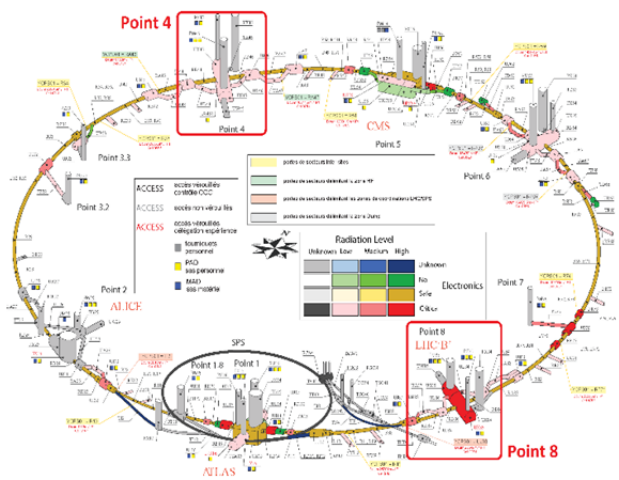


Figure 1: LHC critical areas considered by the R2E Mitigation Project.

A project called Radiation to Electronics (R2E) has been created in 2009 to mitigate the risk of failures due to single events upset (SEU) of the equipment installed in the tunnel and shielded areas. In 2011-2012, due to beam parameters increasing, SEUs had a significant

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

BRINGING QUALITY IN THE CONTROLS SOFTWARE DELIVERY PROCESS

Z. Reszela, G. Cuni, C. M. Falcón Torres, D. Fernandez-Carreiras, G. Jover-Mañas, C. Pascual-Izarra, R. Pastor Ortiz, M. Rosanes Siscart, S. Rubio-Manrique, ALBA-CELLS Synchrotron, Cerdanyola del Vallès, Spain

Abstract

The Alba Controls Section (ACS) develops and operates a diverse variety of controls software which is shared within international communities of users and developers. This includes: generic frameworks like Sardana [1] and Taurus [2], numerous Tango [3] device servers and applications where, among others, we can find PyAlarm [4] and Panic [5], and specific experiment procedures and hardware controllers. A study has commenced on how to improve the delivery process of our software from the hands of developers to laboratories, by making this process more reliable, predictable and risk-controlled. Automated unit and acceptance tests combined with continuous integration, have been introduced, providing valuable and fast feedback to the developers. In order to renew and automate our legacy packaging and deployment system we have evaluated modern alternatives. The above practices were brought together into a design of the continuous delivery pipelines which were validated on a set of diverse software. This paper presents this study, its results and a proposal of the cost-effective implementation.

INTRODUCTION

The ACS designed, constructed and maintain the control systems required to run the facility. The core software components of the Alba control systems, like for example Tango – the middleware framework for building distributed control systems, Sardana – the scientific SCADA oriented to the experiment control and Taurus – the GUI library, are fruits of a collaborative effort of many institutes including Alba [6]. Peripheral components, like Tango device servers, Sardana controllers and macros or Taurus GUIs were either developed in-house or reused from the public repositories. Most of them are developed in Python.

Previously, the software development processes were mainly organized around one-person projects that made it difficult to conduct design questioning discussions and limited the knowledge flow. The software testing was neither formalized nor tried to be automated. This made the project transfers between the developers more difficult, for example on a developer leave. Newcomers, without the complete knowledge of the project, could not feel confident when introducing a change in the code without a way of testing it at the unit or at the system level. In some extreme cases this lead to development downtime periods, to abandoned projects or simply to buggy releases. Subversion (SVN) was the standard version control system (VCS) which did not encourage working branch-per-feature mode. Untraceable commit histories were not helping to enter into the project dynamics. The software packaging and deployment were done manually what is neither interesting nor motivating for the engineers. All the above problems and difficulties were not helping in reducing the long lead time – from the scientist request to the successful use of the software in the experiment.

DEVELOPERS COLLABORATION

Almost two years ago the ACS decided to introduce changes in the software development organization. The in-house projects were transformed from the individual to the group-based efforts. Furthermore, in the case of the two core and initially ACS's internal projects Sardana and Taurus (started at Alba in the previous decade) community-driven development and organization models were introduced. This had an impact on the following parts of the software development process.

Code Design

Internally, developers were organized in groups of 4-6 members – the *Scrum* teams [7]. A mixture of senior and junior developers were selected to build each team. The knowledge transfer activities became a second plan and continuous process. Information about the projects, previously restricted to the privileged project owners, quickly equalized among the team. Agile design practices were introduced: avoidance of upfront designs and plans and promotion of iterative and incremental developments. The Scrum activities brought many interesting design discussions that helped to achieve a better quality of the released products.

In parallel, the Sardana and Taurus project development and decision-taking was opened to a community composed mainly by synchrotrons similar to Alba (DESY in Germany, MaxIV in Sweden Solaris in Poland and ESRF in France), as well as by other institutions, companies and individuals (mostly within the Tango collaboration) who are basing their own developments in them. All these entities actively participate in the community activities. The community model requires remote collaboration tools. Sardana and Taurus are hosted on the Souceforge [8] platform and use a number of provided tools (e.g. issue tracker, mailing lists, wikis, etc.) in the code design processes. Discussions about the critical improvements and modifications are organized and formalized around public processes called Sardana Enhancement Proposal (SEP) and Taurus Enhancement

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

LASER –DRIVEN HADRON THERAPY PROJECT

F. Scarlat, Valahia University of Targoviste, Targoviste, Romania
F. Scarlat, A. Scarisoreanu, National Institute for Laser, Plasma and Radiation Physics-INFLPR, Bucharest-Magurele, Romania
N. Verga, University of Medicine and Pharmacy, Carol Davila, Bucharest, Romania
Fl. Scarlat, Bit Solutions, Bucharest, Romania

Abstract

The laser beam (10 PW, 15 fs, 150 J, 10^{23} W/cm^2) generated by APOLLON Laser System, now under construction on Magurele Platform near Bucharest may also be applied in radiotherapy. Starting from this potential application, location of malign tumors in patient may be situated, e.g., superficial (≤5 cm), semi-deep (5-10 cm) and profound (>10-40 cm). This paper presents the main physical parameters of a research project for a therapy based on hadrons controlled by laser, for the treatment of superficial and semi-deep tumors. Energies required for pin-pointing the depth of such tumors are 50-117 MeV for protons and 100–216 MeV/u for carbon ions. Hadron beams with such energies can be generated by the mechanism Radiation Pressure Acceleration (RPA). Besides, the control systems to provide the daily absorbed dose from the direct and indirect ionizing radiation at the level of the malign tumor of 2 Gy in 1 or 2 minutes with expanded uncertainty of 3 % are presented.

INTRODUCTION

APOLLON Laser System of 10 PW (150 J, 10^{23} W/cm^2, 15 fs) [1], which is under construction on Magurele Platform close to Bucharest, could employ the relativistic laser beams ($I_0\lambda_0^2>10^{18}$ Wcm^{-2}μm^2) as an optional application for researches and experiments to accomplish a final project of laser-driven hadron therapy. According to IAEA-TRS 398 standards, the hadrons must have the kinetic energies of 50 to 250 MeV for protons (P) and 100 to 450 MeV for carbon ions (C) [2].

The hadron beams have the in-depth absorbed dose distributions characterized by small relative doses in the entrance area up to the proximity of the practical path end when the dose is rapidly increasing along a very narrow area in the form of a peak called Bragg Peak (Fig. 1). Function of the hadron energy, this peak is occurring at any depth starting from 2.7 cm (50 MeV-P, 100 MeV-C) up to the depth of 38 cm (250 MeV-P, 450 MeV-C), which corresponds to the 1st stage tumors.

In case of tumors of other stages (2nd-4th) the extension of Bragg–Spread Out Bragg Peak (SOBP) is used, with the residual path centred (focused) in the middle of the tumor. At present, the hadron radiotherapy (RT) employs: conventional accelerators of radiofrequency (RF) of NC type, isochronous cyclotron (IBA Protons P: 235 MeV; RF 106 MHz; SC isochronous cyclotron (Varian P; 150 MeV, 72.8 MHz), SC Syncyclotron (Mevion P: 250, 20t), and Synchrotron: slow-cycling (Siemens C:85-430 MeV/u) and proton linacs [3].

At present there are R&D dedicated to obtaining a compact unit. In view of that, a FFAG is in process to be finalized, a cyclinac (cyclotron + high frequency linac) and a Dielectric Wall Accelerator [4].

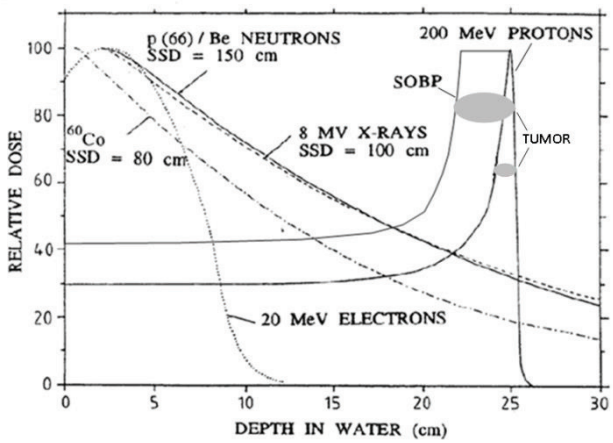


Figure 1: Depth dose distributions in water for X & γ rays, neutron and proton.

Besides, the gradient of the acceleration, limited to 50-100 MV/m in RF linac, may be removed by substituting the RF wave with the laser EM wave that allows an acceleration gradient of 50-100 GV/m [5]. This paper presents the physical parameters of a project employing a laser accelerator for RT of tumors located up to 10 cm in depth using RPA mechanism.

RPA ACCELERATION METHOD

The mechanism called Radiation Pressure Acceleration (RPA) is based on the transfer of energy and the flux of the moment between the linear polarized (LP) laser pulse or circular polarized (CP) laser pulse and the particles inside a solid target. When the target is thin, the acceleration mechanism is called "light sail" (LS) RPA, and when the target is thick (as the case in this paper), the mechanism is called "hole boring" (HB) RPA. The RPA theory is described, for example, in [6].

The HB RPA mechanism developed by Robinson & co. in relativistic regime CP gives the expression for the momentum balance of the plasma surface [7]

$$(2I_0)\left(1/\gamma_f^2\right)\left(1+\beta_f\right)^2 = 2\gamma_f^2 m_i n_i v_f^2 = (2A/Z)\gamma_f^2 m_p n_{e0} v_f^2 \quad (1)$$

where $I_0=(1/2)\varepsilon_0 c E_o^2$ is the laser radiation intensity ($I_0\lambda_0^2$[W/cm^{-2}μm^2]$=1.37\cdot10^{18}\cdot\alpha\cdot a_o^2$ cu α=½ for LP/CP laser pulse), $\varepsilon_o=8.854\cdot10^{-12}$ [As/Vm] is the permittivity of

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

A UNIFIED APPROACH TO THE DESIGN OF ORBIT FEEDBACK WITH FAST AND SLOW CORRECTORS

S. Gayadeen, M. T. Heron, G. Rehm, Diamond Light Source, Didcot, UK

Abstract

A unified control design is proposed to simultaneously determine the inputs to both the fast and slow arrays of correctors. By determining the interaction of the spatial subspaces of each array of correctors, spatial modes which require both fast and slow correctors can be identified. For these modes, a mid-ranging control technique is proposed to systematically allocate control action for each corrector. The mid-ranging control technique exploits the different dynamic characteristics of the correctors to ensure that the two arrays of actuators work together and avoid saturation of the fast correctors. Simulation results for the Diamond Storage Ring are presented.

INTRODUCTION

In this paper, an approach to electron orbit controller design is presented where there are two arrays of actuators with different dynamics. The aim is to design a controller which uses both corrector arrays to meet the electron beam stability requirements so that: the time taken for computation of the control action is not greater than that of the single array system, the tuning for performance is intuitive and the specific dynamics and constraints of each array are addressed. In order to meet these goals, a method of exploiting the knowledge of the spatial responses of the arrays of correctors is used to determine the interaction between the controllable subspaces. As a result the problem can be decomposed into a series of single-input, single-output (SISO), two-input, single-output (TISO) and two-input, two-output (TITO) problems. As a consequence, the online computation is minimised by selecting appropriate control directions when subspaces of the two arrays overlap. In particular, if the control directions of the two arrays are orthogonal, then the problem can be interpreted as a decoupled structure (i.e. two SISO structures) and SISO design techniques can be applied. If however, the control directions align, the problem has a TISO structure, and mid-ranging control is proposed. The term mid-ranging control typically refers to the class of control problems where two actuators are manipulated to control one measured variable. Furthermore there is the condition that one input should return to its midpoint or some setpoint. The inputs usually differ in their dynamic effect on the output and in the relative cost of manipulating each one, with the fast input normally being more costly to use than the slow input [1]. Therefore mid-ranging control schemes seek to manipulate both inputs upon an upset but then gradually reset or mid-range the fast input to its desired setpoint. Mid-ranging control therefore is suitable for the fast and strong corrector problem. Though it may be possible to manipulate each actuator separately, for electron beam control it is desirable to simultaneously manipulate both inputs, as the strong correctors are bandwidth limited. Additionally, an important characteristic of mid-ranging applications is that input constraints on the faster input are avoided. There are several approaches to designing mid-ranging controllers [1], but in [2] an IMC structure is used, which is adopted in this paper since it is consistent with the design for a single array of actuators at Diamond Light Source [3].

MULTI-ARRAY CONTROLLER DESIGN

Subspace Interaction

For a storage ring electron orbit control system with two arrays of actuators where N_f and N_s are the number of actuators in each array, the dynamics of each array $g_{s,f}(z)$ are given by a first order plus delay transfer function determined by the open loop bandwidth $a_{s,f}$ and delay $\tau_{s,f}$ [4], and without loss of generality, it is assumed that:

1. $g_s(1) = 1$ and $g_f(1) = 1$ i.e. the DC gains of the dynamics can be taken as unity.
2. the dynamics of $g_f(z)$ are faster than $g_s(z)$.
3. $N_s \leq N_f$.
4. all the actuators of a given array have the same dynamics.

The position measured at M BPMs is described by

$$Y(z) = g_s(z)R_sU_s(z) + g_f(z)R_fU_f(z) + D(z) \qquad (1)$$

where $U_s(z)$ and $U_f(z)$ represent the inputs applied to the two distinct arrays of actuators and $D(z)$ represent the disturbances acting on the electron beam. The response matrices for each array is represented by R_s and R_f where $N_s = \text{rank}(R_s)$ and $N_f = \text{rank}(R_f)$. The response matrices can each be expressed in terms of reduced singular value decompositions so that,

$$R_s = \Phi_s\Sigma_s\Psi_s^{\mathrm{T}}, \qquad R_f = \Phi_f\Sigma_f\Psi_f^{\mathrm{T}} \qquad (2)$$

where $\Phi_s \in \mathbb{R}^{M\times N_s}$, $\Sigma_s \in \mathbb{R}^{N_s\times N_s}$, $\Psi_s \in \mathbb{R}^{N_s\times N_s}$, $\Phi_f \in \mathbb{R}^{M\times N_f}$, $\Sigma_f \in \mathbb{R}^{N_f\times N_f}$ and $\Psi_f \in \mathbb{R}^{N_f\times N_f}$. The columns of Ψ_s and Ψ_f in Eq. 2 represent the controllable subspaces of the response matrices and even though R_s and R_f are independent, there may be some overlap between the controllable subspaces and the relationship between the subspaces of R_s and R_f can be found by comparing the subspaces of Φ_s and Φ_f which are both orthogonal and can be determined from Algorithm 12.4.3 in [5] such that

$$\Phi_f^{\mathrm{T}}\Phi_s = A \qquad (3)$$

and the singular value decomposition of A is given by

$$A = \Phi_A\Sigma_A\Psi_A^{\mathrm{T}} \qquad (4)$$

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

CONTROL SYSTEM CHALLENGES FROM AN UPGRADE TO THE DIAMOND LIGHT SOURCE STORAGE RING

M. T. Heron, A. Rose, Presented on behalf of the DDBA control and instrumentation team,
Diamond Light Source, Oxfordshire, UK.

Abstract

In 2016 Diamond Light Source will replace one double bend achromatic cell of the storage ring (SR) with two double bend achromatic cells in the same longitudinal space. This will create an additional straight section for an insertion device (ID), thereby converting a bending magnet source point into an ID source point. Installation of the two new cells and recommissioning of the SR will take place in an eight week shutdown, with an expectation that post-shutdown operation will resume with a level of performance comparable to that pre-installation. The additional components in the two new cells necessitate a substantial reworking of the interface layer of control system, together with changes to all applications which are dependent on the physics parameters of the storage ring. This paper will describe how the control system aspects of this project will be managed.

INTRODUCTION

Diamond Light Source is part way through a project to replace one of the standard Double Bend Achromat (DBA) cells of the 561 metre circumference electron storage ring with a new design consisting two DBA cells and a straight section for an insertion device[1]. This arrangement of the two DBA cells has been termed the "Double-Double Bend Achromat" (DDBA) cell. The new insertion device associated with the DDBA cell will service a new photon beamline which otherwise would have had a bending magnet as its source.

The project management aspects of the DDBA project, from a control and instrumentation point of view, are of relevance to future projects at Diamond and other light sources. In the case of Diamond there are presently accelerator physics studies exploring whether a second DDBA can been installed in 2018, together with a conceptual design study considering how the whole storage ring could be rebuilt with a more modern lattice design to reduce the natural emittance. For other light sources, a number of facilities have plans to rebuild their storage rings with more modern lattice designs, again to reduce the natural emittance. As operational facilities, all these projects face the same challenge as the DDBA project in how to make such an invasive change to an operational facility while minimising the impact on user time and maintaining operational performance.

DDBA PROJECT

The existing DBA cell consists of two dipoles and ten quadrupole and seven sextupole magnets to bend and focus the electron beam, which are mounted on three girders. With the DDBA design, each half of the DDBA cell is mounted on one girder including five quadrupoles, five sextupoles and the two dipole magnets. Between the two girders is a new straight section that is used to house the new insertion device, as shown in Fig 1.

To fit a DDBA cell into the longitudinal space of a standard DBA cell necessitates higher field gradient magnets, and so requires smaller apertures. It also necessitates a greater number of magnets, including the addition of trim windings on the dipoles, to enable the existing dipole supply to be used for the main ampere turns and to be able to match the integrated field to that of the standard dipoles. While there are adequate numbers of corrector magnets for orbit correction on the sextupoles, a portion are located over copper vessels, which would significantly impact their dynamic performance in the global Fast Orbit Feedback (FOFB). This is addressed through additional separate correctors over stainless steel portions of the vacuum vessel. The ratings of the new quadrupoles are similar to the old, but the new sextupoles have lower impedance, so in both cases the existing power supplies can be used, albeit with a lower operating voltage for the sextupoles. There is one additional beam position monitor, hence one additional Libera Electron. For the vacuum vessel the use of lumped NEG pumps means fewer ion pumps, but a comparable number of vacuum gauges and additional vacuum valves. For vessel thermal protection there are comparable additional thermal monitors, thermal interlocks and water flows but in different locations. These differences in instrumentation are summarised in Table 1.

ELECTRICAL AND INSTRUMENTATION

From Table 1, the numbers of each component type in the DDBA cell are not significantly different from a standard DBA cell. However, for the majority of electrical circuits it was concluded that it would be better (quicker and operationally more reliable) if new cables were installed to connect the girder back to the instrumentation located outside the ring in the Control Instrumentation Area (CIA). This largely came about because one of the three cable penetrations would not be available after the changes, meaning that a third of the cables would need rerouting, and it was not evident there was sufficient cable length to reassign cables to their new locations on the DDBA cell. Despite the small overall change in numbers of components to control, the new instrumentation necessitated additional racks, which exceed the capacity of the CIA and so necessitate it being extended.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

ROBUST STABILITY ANALYSIS OF ORBIT FEEDBACK CONTROLLERS

S. Gayadeen, M. T. Heron, G. Rehm, Diamond Light Source, Didcot, UK

Abstract

Closed loop stability of electron orbit feedback controllers is affected by mismatches between the accelerator model and the real machine. In this paper, the small gain theorem is used to express analytical criteria for closed loop stability in the presence of spatial uncertainty. It is also demonstrated how the structure of the uncertainty models affects the conservativeness of the robust stability results. The robust stability criteria are applied to the Diamond electron orbit controller and bounds on the allowable size of spatial uncertainties which guarantee closed loop stability is determined.

INTRODUCTION

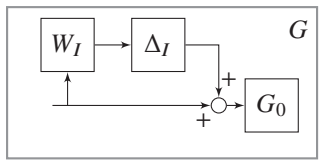
For electron orbit control, the nominal or "golden" response matrix is used to design the controller gains. The control system is referred to as "robust" if it is insensitive to differences between the actual system and the model of the system which is used to design the controller [1]. These differences are referred to as model-plant mismatch or simply uncertainty. The general approach for robust control design is to find a representation of the model uncertainty and then check for robust stability i.e. determine whether the system remains stable for all "real" processes within the uncertainty set. The following notation is adopted: Π is a set of possible perturbed plant models, $G_0(z)$ is the nominal plant model which belongs to the set Π and $G(z)$ is the perturbed model (representing the real plant) which also belongs to the set Π. A norm-bounded uncertainty description is used to describe the uncertainty i.e. the set Π is generated by allowing H_∞ norm bounded perturbations in the nominal plant i.e. Δ is a normalised perturbation with H_∞ norm ≤ 1 and the H_∞-norm of a discrete operator $M(z^{-1})$ is defined as

$$\|M\|_\infty := \max_{\omega\in[-\pi,\pi]} \sigma_{max}(G(e^{j\omega})) \qquad (1)$$

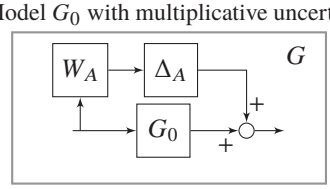
where M is linear, time-invariant and $\sigma_{max}(\cdot)$ indicates the largest singular value of the matrix.

Uncertainty can be classed as either parametric or unmodelled dynamics. Parametric uncertainty is where the model structure is known but some parameters are unknown or uncertain, whereas unmodelled dynamics refers to the case where dynamics have been neglected or missing. In this paper, "lumped" uncertainty is considered which represents one of several sources of parametric uncertainty and/or unmodelled dynamics. This type of uncertainty is usually represented as multiplicative uncertainty [1] as shown in Fig. 1a. In the case of multiplicative uncertainty,

$$\Pi_I : G(z) = G_0(z)(1+W_I(z)\Delta_I(z)); \quad \|\Delta\|_\infty \leq 1 \forall\omega \quad (2)$$



(a) Model G_0 with multiplicative uncertainty.



(b) Model G_0 with additive uncertainty.

Figure 1: Uncertainty representations of real process G.

where the subscript I indicates "input" uncertainty. In the case of additive uncertainty,

$$\Pi_A : G(z) = G_0(z) + W_A(z)\Delta_A(z); \quad \|\Delta\|_\infty \leq 1 \forall\omega \quad (3)$$

which is shown in Fig. 1b. In each case, a weight $W(z)$ is introduced in order to normalise the perturbation to be less than 1 in magnitude at each frequency and to obtain the weight:

- for multiplicative uncertainty:

$$|W_I(e^{j\omega})| \geq \max\left|\frac{G(e^{j\omega}) - G_0(e^{j\omega})}{G_0(e^{j\omega})}\right| \quad \forall\omega \qquad (4)$$

- for additive uncertainty:

$$|W_A(e^{j\omega})| \geq \max|G(e^{j\omega}) - G_0(e^{j\omega})| \quad \forall\omega. \qquad (5)$$

The procedure for robust stability test is as follows:

1. The closed loop system is represented in the structure shown in Fig. 2 where Δ represents the uncertainty in the system and M is the transfer function matrix describing the closed loop "as seen by" Δ.

2. The **small gain theorem** for systems represented as Fig. 2, is applied. The theorem takes the form [1]:

   Theorem 1 (Robust stability) *Assuming that the nominal system $M(e^{j\omega})$ is stable and that the perturbations Δ are stable, then the interconnected $M-\Delta$ system in Fig. 2 is stable for all perturbations Δ, satisfying $\|\Delta\|_\infty \leq 1$ (i.e. the system is robustly stable) if and only if*

$$\|M\|_\infty < 1. \qquad (6)$$

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

UNCERTAINTY MODELLING OF RESPONSE MATRIX

S. Gayadeen, M. T. Heron, G. Rehm, Diamond Light Source, Didcot, UK

Abstract

Electron orbit feedback controllers are based on the inversion of the response matrix of the storage ring and as a result, mismatches between the accelerator model and the real machine can limit controller performance or cause the controller to become unstable. In order to perform stability analysis tests of the controller, accurate uncertainty descriptions are required. In this paper, BPM scaling errors, actuator scaling errors and drifts in tune are considered as the main sources of spatial uncertainties and because most electron orbit feedback systems use Singular Value Decomposition (SVD) to decouple the inputs and outputs of the system, the uncertainty can be expressed in terms of this decomposition. However SVD does not allow the main sources of uncertainty to be decoupled so instead, a Fourier-based decomposition of the response matrix is used to decouple and model the uncertainties. In this paper, both Fourier and SVD uncertainty modelling methods are applied to the Diamond Light Source storage ring and compared.

INTRODUCTION

The response matrix, $R \in \mathbb{R}^{M\times N}$ is the steady state (DC) response of the correctors to a change in the transverse orbit position measured at the mth beam position monitors (BPMs) due to a transverse kick at the nth corrector. For correction of the electron orbit, a map from BPM to corrector is required and therefore the inverse of the nominal (or "golden") response matrix is required. The inverse is normally computed using a Singular Value Decomposition (SVD) approach, such that

$$\begin{aligned} R_o &= \Phi_o \Sigma_o \Psi_o^{\mathrm{T}} \\ R_o^{-1} &= \Psi_o \Sigma_o^{-1} \Phi_o^{\mathrm{T}}. \end{aligned} \tag{1}$$

However because of BPM and corrector scaling errors from imperfect calibrations or drifts in the tune, the actual response of the machine may differ over time from the ideal response. This means that the control action calculated from the ideal response may be either too aggressive or too conservative. In the case of such model mismatch, applying the control action may result in closed loop instabilities or even if the closed loop remains stable, steady state performance will degrade. By modelling the various sources of the uncertainties, the mismatch between the machine response and the ideal response can be quantified and the uncertainty description can then be used in stability analysis techniques to determine the allowable bound on the size of uncertainty to ensure closed loop stability ([1] in these proceedings). In this paper, modelling of the uncertainties associated with the response matrix is presented. Firstly, the uncertainties are described within the commonly used SVD framework. However there are limitations to this approach and as an alternative, a harmonic decomposition of the response matrix is presented and the uncertainty is expressed in terms of the Fourier coefficients of the response.

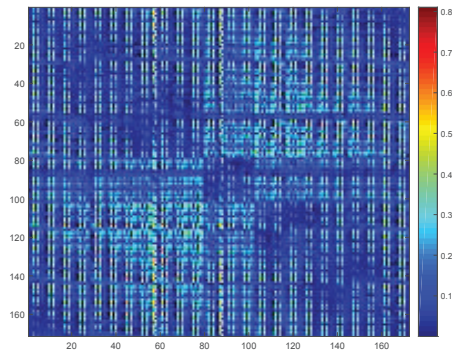


Figure 1: Additive uncertainty in response matrix R_0 ($|\Delta_R|$).

UNCERTAINTY DESCRIPTIONS USING SINGULAR VALUE DECOMPOSITION

The uncertainty in the response matrix can be expressed by an additive error,

$$R = R_o + \Delta_R \tag{2}$$

where R represents a measured response matrix and R_o represents the nominal response matrix. Figure 1 shows the structure of Δ_R for the Diamond storage ring taken from one measurement of R. The uncertainty in the response matrix combines the errors in Φ, Σ and Ψ and it is not straightforward to distinguish, for example, the effect of BPM errors from corrector errors. It is therefore more useful to model the uncertainty in the process output (sensor values) by multiplicative operators [2], which can be expressed as

$$R = (I + \Delta_B) R_0 \tag{3}$$

so that

$$\Delta_B = R R_0^{-1} - I. \tag{4}$$

The uncertainty can be projected into "mode space", such that

$$\Delta_B = \Phi \Sigma \Psi^{\mathrm{T}} \Psi_0 \Sigma_0^{-1} \Phi_0^{\mathrm{T}} - I. \tag{5}$$

Likewise, the same treatment can be given to uncertainties in the process input (corrector values), which is expressed as

$$R = R_0 (I + \Delta_C) \tag{6}$$

so that

$$\Delta_C = \Psi_0 \Sigma_0^{-1} \Phi_0^{\mathrm{T}} \Phi \Sigma \Psi^{\mathrm{T}} - I. \tag{7}$$

By assuming that all the uncertainty lies in either Φ or Ψ, then the uncertainties can be expressed as

$$\begin{aligned} \bar{\Delta}_B &= \Phi \Phi_0^{\mathrm{T}} - I \\ \bar{\Delta}_C &= \Psi_0 \Psi^{\mathrm{T}} - I. \end{aligned} \tag{8}$$

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

STATUS OF THE SOLARIS CONTROL SYSTEM - COLLABORATIONS AND TECHNOLOGY *

P. Goryl#, C.J. Bocchetta, P. Bulira, Ł. Dudek, P. Gałuszka, A. Kisel, W. Kitka, M. Kopeć, P. Kurdziel, M. Ostoja-Gajewski, M.J. Stankiewicz, J. Szota, A.I. Wawrzyniak, K. Wawrzyniak, Ł. Żytniak, Solaris, Krakow, Poland
T. Szymocha, ACK Cyfronet AGH, Krakow, Poland
V. Hardion, J. Jamróz, J. Lidon-Simon, M. Lindberg, A.G. Persson, D. Spruce, MAX IV Laboratory, Lund, Sweden
I. Dolinsek, Cosylab, Ljubljana, Slovenia

Abstract

The Solaris is a synchrotron light source starting just now in Krakow, Poland. It is built with strong collaboration with other European accelerator facilities. The MAX-IV project in Lund, Sweden and Tango Community are the most important partners in the project in respect to the control system. Solaris has built a twin copy of MAX-IV 1.5GeV ring and linear accelerator based on the same components as the ones of MAX-IV. Thus, both facilities share know-how and apply similar technologies for the control system, among them the Tango CS is used for software layer. Status of the control system in Krakow as well as collaborations and technological choices impact on its success are presented in this paper.

THE ACCELERATOR AND BEAMLINES

The Solaris machine is 1.5GeV storage ring the same as one of the MAX-IV supplied with a linear accelerator providing electron beam up to 600MeV. Since the linac is not providing full energy beam there is need of energy ramping in the storage ring. This is a difference to MAX-IV setup. The Solaris project includes also two beamlines. One is using bending magnet radiation with XAS and PEEM end-stations and the other is basing on undulator source with an UARPES end-station [1].

The accelerators and beamlines are installed except few components of PEEM beamline and two Landau cavities in the storage ring. There is on-going commissioning of the accelerator [2]. The latest result is accumulation of 20mA current in the storage ring which then has been ramped to full energy of 1.5GeV. Currently certain optimization has been done that let us by the day of writing this paper (17.10.2015) accumulate the current of 48mA. The work is in progress to do energy ramping of this much od current. For now limitation are instabilities due to the ion trapping process.

There is planned upgrade shutdown starting end of November 2015 till beginning of January 2016. During this period the landau cavities will be installed as well as missing components of the PEEM beamline. Network infrastructure will be supplemented with a new core switch to enable 10GBps connections through the whole network. It is also planned to do upgrade of the Tango control system to the newest version of Tango 9.1.

After the shutdown commissioning of the beamlines will proceed. It is expected that the Solaris will be ready for the first external users in second half of 2016.

*Work supported by the European Regional Development Fund within the frame of the Innovative Economy Operational Program: POIG.02.01.00-12-213/09
#piotr.goryl@uj.edu.pl

THE CONTROL SYSTEM

The control system of the accelerators and beamlines has been deployed and is operational. All major issues has been debugged and solved. However, few low level components still need optimization. Certain high level applications for future simple operation are upon development.

Tango Control System

The Tango Controls [3] is used for systems integration. Since the Solaris and MAX-IV have similar systems most of the software components are shared and lot of them have been developed in Lund [4]. Solaris and MAX-IV are developing their applications with the Python programming language. The PyTango, the Taurus and the Sardana packages developed at ALBA are used for this purpose. An archiving system which role is to store history of controlled signals is using HDB and TDB tools developed at Soleil [5].

Some of the device servers provided by the Tango Community [6] were reused.

The integration work has been contracted by Solaris to a commercial company Cosylab (Ljubljana, Slovenia). They have also developed few device servers and part of GUI applications including an open source ControlProgram which is an application being an entry point to all control system functionalities needed for operation [7]. The ControlProgram has been published on the public Solaris GitHub repository [8].

The Matlab Midlayer with Tango-Matlab binding [9] is used for machine physics control. The basic configuration has been done by the Cosylab then the work has been overtaken by Solaris. The work were conducted with support from Soleil.

Control System Hardware

The timing system has been provided within the contract with Cosylab. The system based on Micro Research

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

DETECTOR CONTROLS MEETS JEE ON THE WEB

F. Glege, M. Janulis, A. Andronidis, O. Chaze, C. Deldicque, M. Dobson,
A. Dupont, D. Gigi, J. Hegeman, R. Jiménez Estupiñán, L. Masetti, F. Meijers,
E. Meschi, S. Morovic, C. Nunez-Barranco-Fernandez, L. Orsini, A. Petrucci, A. Racz,
P. Roberts, H. Sakulin, C. Schwick, B. Stieger, S. Zaza, P. Zejdl (CERN, Geneva, Switzerland),
U. Behrens (DESY, Hamburg, Germany), O. Holme (ETH Zurich, Switzerland),
J. Andre, R. K. Mommsen, V. O'Dell (Fermilab, Batavia, Illinois, USA), G. Darlea,
G. Gomez-Ceballos, C. Paus, K. Sumorok, J. Veverka (MIT, Cambridge, Massachusetts, USA),
S. Erhan (UCLA, Los Angeles, California, USA),
J. Branson, S. Cittolin, A. Holzner, M. Pieri (UCSD, La Jolla, California, USA)

Abstract

Remote monitoring and controls has always been an important aspect of physics detector controls since it was available. Due to the complexity of the systems, the 24/7 running requirements and limited human resources, remote access to perform interventions is essential.
The amount of data to visualize, the required visualization types and cybersecurity standards demand a professional, complete solution.
Using the example of the integration of the CMS detector controls system into our ORACLE WebCenter infrastructure, the mechanisms and tools available for integration with controls systems shall be discussed. Authentication has been delegated to WebCenter and authorization been shared between web server and control system. Session handling exists in either system and has to be matched. Concurrent access by multiple users has to be handled. The underlying JEE infrastructure is specialized in visualization and information sharing. On the other hand, the structure of a JEE system resembles a distributed controls system. Therefore an outlook shall be given on tasks which could be covered by the web servers rather than the controls system.

INTRODUCTION

The Compact Muon Solenoid (CMS) detector is one of the four experiments of the Large Hadron Collider (LHC) at CERN. The Detector Controls System (DCS) is needed to bring the detector in a state to collect physics data. The scale of the system is illustrated with the following numbers:

- ~3 million parameters
- ~700.000 lines of code
- ~35000 finite state machine nodes
- 34 SCADA systems
- 29 redundant PC pairs (Windows)
- ~50 DB schemas (ORACLE)
- O(TB) of data in database schemas

The system is running in production since several years with a high efficiency. It is built using the SCADA system WinCC OA from the Austrian company ETM. At CERN the Joint Controls Project (JCOP) created a toolkit based on WinCC OA to facilitate the development of the experiment controls systems. Today this is used by all LHC experiments and for several other projects at CERN.

PART1: REMOTE CONTROL VIA WEB

Visualization in WinCC OA

In WinCC OA visualization is done through so-called panels. Using a graphical editor, graphics objects (buttons, text, lines, etc.) can be placed on a window. Each element can be enhanced by scripts, which allow to change the properties of the element itself. The scripts have access to the process data and can therefore reflect the process information through the visualization. The graphical editor produces a file describing the panel layout and functions. This file is then used by a WinCC OA process to visualize these panels. This process uses Qt as a graphics library, with all graphics elements being drawn using the Qt graphics API. The Qt Library itself translates the drawing of the panel into primitive paint commands which are sent to the device used for visualization.

Remote Control

Around 4000 members of the CMS collaboration should be able to access the DCS information remotely in a read only manner. The data should ideally be live to ensure a good user experience. Experts should in addition be able to interact and remotely control their systems. Authorization and authentication are required to enable interaction.
The means for remote control currently used in CMS are

- snapshots taken regularly of selected panels being served as images vie the web server
- HTML/java script pages providing live information collected
  - from the logging database
  - through protocols supported and used by the controls system

A full discussion of the current remote control solution in CMS DCS can be found in [1].

Live Web Access

Thousands of panels have been created for the detector controls of the LHC experiments. Reimplementing all of them at once to provide a pure HTML solution which

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

CS-STUDIO SCAN SYSTEM PARALLELIZATION*

K.U. Kasemir, M.R. Pearson, ORNL, Oak Ridge, TN37831, USA

Abstract

For several years, the Control System Studio (CS-Studio) Scan System has successfully automated the operation of beam lines at the Oak Ridge National Laboratory (ORNL) High Flux Isotope Reactor (HFIR) and Spallation Neutron Source (SNS). As it is applied to additional beam lines, we need to support simultaneous adjustments of temperatures or motor positions.

While this can be implemented via virtual motors or similar logic inside the Experimental Physics and Industrial Control System (EPICS) Input/Output Controllers (IOCs), doing so requires a priori knowledge of experimenters requirements. By adding support for the parallel control of multiple process variables (PVs) to the Scan System, we can better support ad hoc automation of experiments that benefit from such simultaneous PV adjustments.

MOTIVATION

EPICS has been used with great success on the SNS accelerator, and is now a key component of on-going beam line software updates [1]. Compared to the accelerator, each beam line presents a much smaller number of devices to control. At the same time, the beam line environment is more flexible as devices are added, replaced, or operated in different ways.

The CS-Studio Scan System was developed to allow the flexible assembly and execution of "recipes", that is lists of commands [2]. After gaining operational experience on a couple of beam lines, the need for parallel command execution became obvious.

For example, the control of incident beam energy typically requires the adjustment of several neutron chopper settings as well as motor positions for slits or sample orientation. Executing these in parallel instead of sequentially allows for a significantly faster experiment preparation.

Another example is the combined movement of motors or temperatures for ad-hoc experiments. Ideally, the need for the parallel movement of motors is known in advance. Choosing suitable motor controller hardware will then allow configuring the desired motion curves into the hardware, resulting in the most efficient and accurate position control. In reality, the need for ganged motor movement can arise without prior notice. The control system software then needs to make a reasonable attempt at combined motor movement.

*This manuscript has been authored by UT-Battelle, LLC under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (http://energy.gov/downloads/doe-public-access-plan).

SCAN SYSTEM

Basic Scan Commands

While the scan system supports several commands [2], it is most important to understand the fundamental "Set" command, which writes a value to an EPICS channel:

```
Set("MotorChannel", value=14)
```

It can wait for a read-back to match the written value, with configurable tolerance and timeout in seconds:

```
Set("MotorChannel", value=14,
    readback="Motor.RBV", tolerance=0.1,
    timeout=30)
```

Alternatively, it can await "completion":

```
Set("MotorChannel", value=14,
    completion=True,
    timeout=30)
```

Finally, it can await "completion" and then check a read-back for matching the written value:

```
Set("MotorChannel", value=14,
    completion=True,
    readback="Motor.RBV", tolerance=0.1,
    timeout=30)
```

In the last case, the original implementation used the timeout both to await completion and then to wait for the read-back to match. Operationally, we found it more practical to only apply the timeout to the completion. Once the channel completes, the read-back is expected to match right away as soon as the scan server reads the most recent value of the read-back channel.

A motor-related channel will signal "completion" after the controller hardware moved the motor toward the commanded position, verified the encoder read-back, maybe performed several retries, corrected for backlash, and finally confirmed that the motor rests at the requested position.

If a device does not support "completion", it is tempting to rely on only the read-back as in the second example given above. That command will indeed wait until the read-back of the device matches the desired value, but it is misleading. The read-back may temporarily match during an overshoot or backlash

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

ADVANCED WORKFLOW FOR EXPERIMENTAL CONTROL

D. Mannicke, N. Hauser, N. Xiong, ANSTO, Sydney, Australia

Abstract

GumTree is a java software product developed at ANSTO and used for experimental control as well as data visualization and treatment. In order to simplify the interaction with instruments and optimize the available time for scientists, a user friendly multi sample workflow has been developed for GumTree. Within this workflow users follow a step by step guide where they list available samples, setup instrument configurations and even specify sample environments. Users are then able to monitor the acquisition process in real-time and receive estimations about the completion time. In addition users can modify the previously entered information, even after the acquisition sequence has commenced. This paper will focus on how ANSTO integrated a multi sample workflow into GumTree [1], what approaches were taken to allow realistic time estimations, what programming patterns were used to separate the user interface from the execution of the acquisition, and the future opportunities the approach enables.

INTRODUCTION

Experiments have become more complex and involve different instrument configurations and sample environments. Our aim was to provide scientists with a user friendly application that would simplify the data acquisition process and the associated interactions with the instrument being used as well as help them to optimise the experimental time available.

For example with Quokka, a small angle neutron scattering beamline at ANSTO, it is very common to set up an experiment that requires measurements with 3 configurations, at least 7 samples and may also involve 5 different temperatures. This setup alone results in 105 single measurements which would be very time consuming to manually prepare.

In order to increase the acceptance of the provided solution by the scientific community, scientists and researchers were included in the development process. We asked about their expectations, ideas and experiences at other facilities. We received very different and sometimes contradicting responses. Finally we decided on a solution that delivered a compromise. The end result is suitable for the majority of experiments conducted on the beamlines.

Initially we assumed it would be sufficient to simply loop-over configurations, samples and temperatures in order to generate an acquisition sequence using existing scripting language commands. However, in practice many scientists desired more flexibility; they wanted to be able to adjust the order of the acquisition sequence including duplicate and remove measurements. The scientists also wanted the ability to make these changes after the acquisition sequence had commenced. The interactivity on a running command sequence could neither be implemented with a scripting language 'loop' construct nor with a simple 'command list' or 'command table', which are traditionally generated by unrolling nested loops.

A key requirement from the development perspective is that the system could be deployed and maintained on multiple beamlines with minimal effort. Therefore we had to develop a solution that keeps most of the implementation generic but provides sufficient flexibility to allow a customised user interface to be created for each beamline.

In addition, to improve the robustness it was decided to separate the system into a server and a rich client application. The server was developed to run in a headless mode on hardware with direct communication to the instrument control system. Furthermore, the server had to provide an interface that allows multiple clients to monitor and control the acquisition sequence.

IMPLEMENTATION

For the implementation of the user interface, the workflow was divided into four individual tabs where each tab represents a different aspect of the system (See Fig. 1). The first tab is dedicated to the sample stage; this includes specifying the properties of the used samples as well as the utilized sample holder. The second tab is concerned with setting up the desired beamline configurations which may include different sample to detector distances and aperture configurations. The next tab is concerned with sample environments; this is where a user can specify different temperature ramps or magnetic field strengths. On the final tab, the user is able to start and stop the acquisition as well as control and monitor the acquisition sequence.

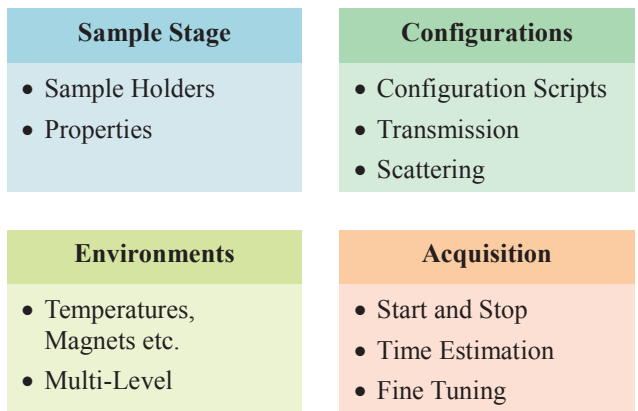


Figure 1: Workflow aspects

In order to keep the system easy to maintain for multiple beamlines we minimized the instrument specific components (See Fig. 2). For example on the server side, only an XML Schema Definition file, which is used as an

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

ITERATIVE DEVELOPMENT OF THE GENERIC CONTINUOUS SCANS IN SARDANA

Z. Reszela, G. Cuni, C. M. Falcón Torres, D. Fernandez-Carreiras, C. Pascual-Izarra, M. Rosanes Siscart, ALBA-CELLS Synchrotron, Cerdanyola del Vallès, Spain

Abstract

Sardana [1] is a software suite for Supervision, Control and Data Acquisition in scientific installations. It aims to reduce cost and time of design, development and support of the control and data acquisition systems. Sardana is used in several synchrotrons where continuous scans are the desired way of executing experiments [2]. Most experiments require an extensive and coordinated control of many aspects like positioning, data acquisition, synchronization and storage. Many successful ad-hoc solutions have already been developed, however they lack generalization and are hard to maintain or reuse. Sardana, thanks to the Taurus [3] based applications, allows the users to configure and control the scan experiments. The MacroServer, a flexible Python [4] based sequencer, provides parametrizable turn-key scan procedures. Thanks to the Device Pool controllers interfaces, heterogeneous hardware can be easily plug into Sardana and their elements used during scans and data acquisitions. Development of the continuous scans is an ongoing iterative process and its current status is described in this paper.

INTRODUCTION

Traditionally, a continuous scan needed an ad-hoc controls software being in charge of the configuration, process control, data acquisition and storage of the whole measurement. Its use was restricted to particular axes, hardware involved to count encoder or motor pulses and generate external triggers or specific experimental channels and detectors. The generalization of this type of scans is strategic in all synchrotrons and potentially in any other laboratory. A generic continuous scan framework provides the flexibility of step scans in a continuous and synchronized data acquisition, allowing time-resolved experiments. The complexity resides in using any arbitrary combination of movable elements, experimental channels, and detectors, generating the required triggers and evidently, allowing the coexistence of slow interpolated software-triggered channels and “virtual” elements like pseudocounters and pseudomotors.

EXPERIMENT CONTROL WITH SARDANA

Sardana has been designed for suiting large installations like a particle accelerator, or smaller such as experimental stations, up to small labs. Different characteristics and necessities of these type of facilities require from Sardana high flexibility and scalability.

In order to speed up the learning curve many Sardana features were inspired on SPEC [5], a complete and powerful software tool, very popular in building control systems for X-ray and neutron experimental stations. But Sardana goes beyond its functionalities. Its architecture is based on the client-server model with Tango [6] as the middleware. This allows to balance the workload in a distributed manner between multiple clients and multiple servers on both Linux and Windows platform PCs.

Sardana had proven to be a solution for synchrotrons, with its successful use in all beamlines, the accelerators and the peripheral laboratories of the Alba synchrotron [7]. It is also used in other similar installations like DESY in Germany, MaxIV in Sweden Solaris in Poland.

Taurus Based User Interfaces

Users can choose between Taurus based command line (CLI) and graphical interfaces (GUI) for interacting with the Sardana system.

From one hand exist turn-key applications. Spock, a single point of access CLI, provides the total control over the system. It allows to run the control procedures and interact with the laboratory elements. Its syntax tries to mimic as much as possible the SPEC commands, what makes the user transition between the two systems easier. A collection of Sardana specific, Taurus based, widgets (many of them as well available as standalone applications) includes: the experiment configuration tool, the scan plots, the user procedure editor, executors and sequencer as well as the elements' control widgets e.g. a generic motor widget.

From the other hand, users may create their own complete Sardana GUIs without the need to program a single line of code. The TaurusGUI framework [3] offers a wizard-based GUI creation process, in which the user just needs to specify which Sardana system the application should interface to and which graphical synoptic should it use. The application gets automatically populated with the previously listed widgets and the instruments' control panels bidirectionally connected with the graphical synoptic. At the application runtime, user can easily navigate in the laboratory by selecting instruments' panels via clicks on the synoptic.

Standard and Custom Experiment Procedures

Experiment and control procedures in Sardana are programmed as Python scripts and are called macros. The MacroServer manages available to the user macros and executes them either sequentially or simultaneously on the user request. Macro development and execution

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

THE MODULAR CONTROL CONCEPT OF THE NEUTRON SCATTERING EXPERIMENTS AT THE EUROPEAN SPALLATION SOURCE ESS

T. Gahl, R. J. Hall-Wilton, O. Kirstein, T. Korhonen, A. Sandstroem, I. Sutton, ESS Lund, Sweden
T. R. Richter, J. Taylor, ESS Copenhagen, Denmark

Abstract

The European Spallation Source (ESS) in Lund, Sweden has just entered into neutron beam line construction starting detailed design in 2015. As a collaboration of 17 European countries the majority of hardware devices will be provided in-kind. This presents numerous technical and organizational challenges for the construction and the integration of the neutron instrumentation into the facility wide infrastructure; notably the EPICS control network and the facilities absolute timing system.

In this contribution we present a strategy for the modularity of the instruments hardware with well-defined standardized functionality and a minimized number of control & data interfaces. Key point of the strategy is the time stamping of all readings from the instruments control electronics extending the event mode data acquisition from detected neutron events to all metadata. This gives the control software the flexibility necessary to adapt the functionality of the instruments to the demands of each single experimental run. We give examples for the advantages of that approach in classical motion control as well as in complex robotics systems and discuss matching hardware requirements necessary.

INTRODUCTION

The European Spallation Source (ESS) is designed as a long pulse neutron source with the maximum overall flux as its main objective. With a 5MW linear proton accelerator it will be the world's most powerful source of neutrons for science applications. Integrated flux levels will be much higher than existing facilities and the geographical layout will comprise instruments of 160m and longer spanning over 3 instruments halls. At the same time the experimental end stations (instruments) for these kinds of sources need to be more flexible and complex with up to 20 choppers along the beam line. All this presents operational challenges that are best addressed with a good mixture of techniques from existing neutron facilities and other disciplines like x-ray experiments, fundamental physics or industrial applications. Advanced grounding concepts, high rate data acquisition, flexible experimental setup, advanced remote diagnostics tools, pre-emptive maintenance are a few keywords for that.

At the same time the ESS project represents also an organisational challenge with most of the instruments hardware provided in-kind by the 17 European partner countries. Integration of these contributions either as single instruments components or complete instruments into the ESS infrastructure is only possible with a flexible and modular instrument control system. This also supports clear definitions of functionalities and interfaces of the modules. We will present the evolution of our ideas of such a system first formulated in [1] following the current design phases and addressing technical, operational and organisational challenges of the future ESS.

ESS NEUTRON BEAM

Neutron Beam Characteristics

ESS is designed as a long-pulse neutron source with 14Hz repetition rate and 4% duty cycle resulting in 2.86 ms pulse length [2]. With an average power of 5MW of the proton accelerator it will exceed existing neutron sources - either steady state or pulsed - in average neutron flux as well as in peak brightness.

Neutron Beam Development

The energy spectrum of the neutron pulse will range from thermal to cold neutrons with travelling speeds between 2000 and 200 m/s. After a certain travel distance in the neutron guide the pulse will transform into a constant beam modulated in time by neutrons with different energy (Figure 1). Mechanical choppers are used to form a slit system in the time domain to enhance the energy resolution of the resulting neutron beam and/or to limit the spectrum used for the experiments. The exact configuration is determining the length of the instruments ranging from 30 to 160m.

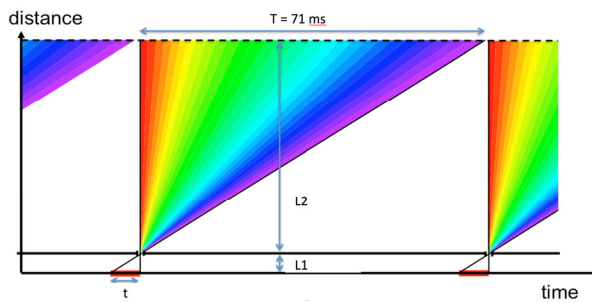


Figure 1: Time-distance diagram of beam development.

Due to the time characteristics of the final neutron beam at the experiments synchronisation and referencing of all experimental data to the proton pulse is a key requirement of instruments control and data acquisition at ESS.

ESS CONTROL INFRASTRUCTURE

EPICS Control System (ICS)

ESS will use EPICS as a global control system for instruments, machine (accelerator, target), central infrastructure and part of conventional facilities. It acts as a

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

THE LMJ SYSTEM SEQUENCES ADAPTABILITY (FRENCH LASER MEGAJOULE)

Y.Tranquille-Marques, J. Fleury, H. Graillot, CEA, CESTA, Le Barp, France
N. Chapron, S. Bailleux, AREVA TA, Le Barp, France
J. Gende, ALTEN, Mérignac, France
J. Nicoloso, CEA, DIF, Bruyères le Châtel, Arpajon, France

Abstract

The French Atomic and Alternative Energies Commission (CEA : Commissariat à l'Energie Atomique et aux Energies Alternatives) is currently building the Laser MegaJoule facility. In 2014, the first 8 beams and the target area were commissioned and the first physics campaign (a set of several shots) was achieved. On the LMJ, each shot requires more or less the same operations except for the settings that change from shot to shot. The supervisory controls provide five semi-automated sequence programs to repeat and schedule actions on devices. Three of them are now regularly used to drive the LMJ. Sequence programs need to have different qualities such as flexibility, contextual adaptability, reliability and repeatability. Currently, the calibration shots sequence drives 328 actions towards local control systems. However, this sequence is already dimensioned to drive 22 bundles, which will lead to manage almost 5300 actions. The presentation introduces the organization of the control system used by sequence programs, the sequence adjustments files, the GRAFCETs of sequences, the GUIs, the software and different tools used to control the facility.

LASER MEGAJOULE SHOT OVERVIEW

The LMJ facility will count 176 laser beams. It is dimensioned to deliver 1.4 MJ shot of UV laser at 351 nm to a 10mm target. The LMJ is designed to study high density plasma physics. It's also a part of the French Simulation Program that forms the basis of the safety and reliability of French nuclear weapons.

The LMJ facility is divided in 4 laser bays. In its center, takes place the target bay. A shot is composed of laser pulses which are amplified in the laser bays and focused on a target in the target bay.

A timing shot that lasts 1s, consists in switching on the laser amplifiers lamps, creating, shaping and pre-amplifying the laser pulses, sending it through the amplification section to fill up, converting laser infrared pulses into laser UV pulses. Then, the energy heats the target and the diagnostics record the experiment radiations data. [1]

Shot sequence manages shot preparation steps, timing shot and post shot step.

First comes facility preparation step. It does not include time critical actions.

The preamplifiers perform several triggers to ensure delivered energy repeatability. If needed, inserters carrying diagnostics move at that time.

Then we have the Alignment step. Each laser beam must hit the millimeter target after hundreds meters propagation. Different areas get their own alignment: Amplification sections, Transport Sections, Frequency converter Sections and Target Chamber.

Once alignment is finished, a whole test without Power conditioning is performed: the validation shot or rod shot. It checks that all the required equipment is activated without wasting the target. This equipment is configured and triggered, laser beam measures are analyzed by automatic data processing.

When the validation shot has passed, power shot repeats the validation shot, adding the power amplifiers charge leading to timing shot.

After the shot, amplifier flash lamps test is performed while other devices are secured. Results are stored in a shot database. This is post shot step.

CONTROL SYSTEM

Data Model: Key Words

The main concept of the data model is called a "resource". A resource represents an equipment (motors, instrumentations, diagnostics…) or a high level function (alignment, amplification section, power conditioning bundle). LMJ counts 200 000 resources.

A resource owns a name and attributes:

- Reserved : to forbid unexpected use,
- Context: resource membership of a perimeter. A perimeter is a group of resources existing for the same purpose such as sequence.
- Operational State : Running, Standby or Stop

A resource can be attached to other objects:

- Functions : what a resource can do,
- Control points : resource information,
- Alarms: asynchronous resource events.

Architecture

The LMJ control system architecture is built on a 4 layers architecture as shown in Fig. 1:

- Level 3 system supervisory provides facility planning and operation functions which are dedicated to the shot director
- Level 2 consists of the supervisory system (SVP) which provides GUI to the shot director. As system sequences are SVP programs, system sequences

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

MACHINE PROTECTION AND INTERLOCK SYSTEMS FOR LARGE RESEARCH INSTRUMENTS

R. Schmidt, CERN, Geneva, Switzerland

Abstract

Major research instruments such as accelerators and fusion reactors operate with large amount of power and energy stored in beams and superconducting magnets. Highly reliable Machine Protection systems are required to operate such instruments without damaging equipment in case of failure. The increased interest in protection is related to the increasing beam power of high-power proton accelerators such as ISIS, SNS, ESS and the PSI cyclotron, to the large energy stored in the beam (in particular for hadron colliders such as LHC) and to the stored energy in magnet systems such as for ITER and LHC. Machine Protection includes process and equipment monitoring, a system to safely stop operation (e.g. dumping the beam or extracting the energy stored in the magnets) and an interlock system for highly reliable communication between protection systems. Depending on the application, the reaction of the protection function to failures must be very fast (for beam protection systems down to some µs). In this presentation an overview of the challenges for protection is given, and examples of interlock systems and their use during operation are presented.

INTRODUCTION

Accelerators, as all other technical systems, must respect some general principles with respect to safety and protection. Protection of people from different threats (radiation, electrical, oxygen deficiency, ..) has always the highest priority and follows legal requirements. The main strategy to protect people during operation is to keep them away from hazards, ensured by a personnel access system. Protection of the environment is the second priority. In this paper the protection of equipment (the investment) is discussed, e.g. accelerators and beam targets, experiments and fusion reactors. Designing a machine protection system is challenging and requires an excellent understanding of the system and its operation, to anticipate and avoid possible failures that could lead to damage or mitigate the consequences.

Protection of accelerators was to topic of a recent Joint Accelerator School, the proceedings with many relevant contributions will be published by the end of 2015 [1].

In general, risks come from energy stored in a system (measured in Joule) as well as from power when operating the system (measured in Watt). Particle accelerators and fusion reactors are examples of such systems, since they operate with large amount of electrical power (from a few to many MW). The energy and power flow needs to be controlled. An uncontrolled release of energy or an uncontrolled power flow can lead to unwanted consequences such as damage or activation of equipment and loss of time for operation.

MACHINE PROTECTION AT LARGE RESEARCH INSTRUMENTS

Many accelerators operate with high power beams or beams with a large amount of stored energy. Accelerators operating with superconducting magnets require sophisticated systems to protect the magnets in case of a quench.

Accelerators with large stored energy in the beams are hadron synchrotrons and colliders (Figure 1), e.g. LHC, RHIC, SPS and FCC (a study for a proton collider with a circumference of 100 km and 100 TeV cm energy, [2]),

Many proton accelerators operate with high power beam, such as the PSI cyclotron, ISIS, SNS, JPARC and ESS (planned to start operation in 2019 [3]). A particular challenge are future ADS machines (Accelerator Driven Spallation) that require to operate with very high beam power and extremely high availability [4]. A prototype is expected for the next decade.

Machine protection is also relevant for some electron accelerators, e.g. free electron lasers (XFEL), synchrotron light sources, e+e- circular and linear colliders. In particular protection at linear colliders is challenging, with MW beam power and tiny beam spot sizes.

Installations with large superconducting magnet systems are also considered, e.g. fusion reactors such as ITER operate with magnets that store more than 50 GJ of energy.

To get an idea what this amount of energy means some examples are given. The energy of pistol bullet is about 500 J, the energy of 1 kg TNT about 4 MJ. The energy of 1 l fuel is about 36 MJ, to melt 1 kg of steel about 800 kJ are required (the energy to melt 1 kg of copper is similar). An accidental release of an energy above one MJ can cause significant damage. With the energy stored in the ITER toroid magnet it is possible to melt 60 tons of copper. Even an accidental release of a small amount of energy in the order of some hundred Joule can lead to some (limited) damage if the energy is released in sensitive equipment.

HAZARDS AND RISKS

Definition of Risk

A hazard is a situation that poses a level of threat to the installation. Hazards are dormant or potential, with only a theoretical risk of damage. Once a hazard becomes "active" it becomes an incident or accident. An accident is defined as an unfortunate incident that happens unexpectedly and unintentionally, typically resulting in damage or injury.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

DESIGN, IMPLEMENTATION AND SETUP OF THE FAST PROTECTION SYSTEM FOR CSNS *

D. P. Jin#, Y. L. Zhang, P. Zhu, IHEP, Beijing, China; DNSC, Dongguan, China

Abstract

The China Spallation Neutron Source (CSNS) [1, 2] is being constructed in Dongguan, Guangdong Province, China. The first proton pulse will hit on the target in March, 2018 as planned. CSNS is a large neutron beam research facility with a proton beam power of 100 kW. Three neutron instruments are being constructed for the first stage, with another 17 to be built later. The Fast Protection System (FPS) is one of the key systems to prevent the accelerator components along the proton beams from being damaged. Design, implementation, setup and tests of the FPS are introduced in this paper.

INTRODUCTION

As the other high power accelerators in the world, components along the proton beam will possibly be damaged by a mis-steering beam, especially for the DTL (Drift Tube LINAC) cells. Another issue is the post irradiation due to the beam loss, which should be controlled strictly below 1 W/m for hand maintenance [3, 4]. A Fast Protection System (FPS) with a response time of around 10 μs is strongly needed. A self-designed FPS has been developed and preliminary tests has been brought out. Till now, most of the signals from the power supplies and beam loss monitors along the linear accelerator are collected by the Fast Protection System.

DESIGN AND IMPLEMENTATION OF THE FPS FOR CSNS

FPGA is chosen as the base of the design since its flexibility and fast response time. High speed serial transmission technique is used to reduce the number of connections tremendously. A 6U standard VME crate is used to ease the maintenance and upgrade, and to improve the reliability and scalability. Figure 1 shows the distribution of the FPS.

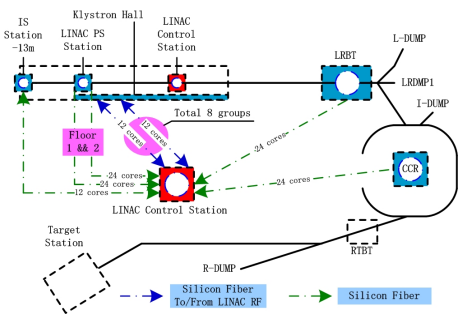


Figure 1: Distribution of the FPS.

Figure 2 shows the structure of the FPS.

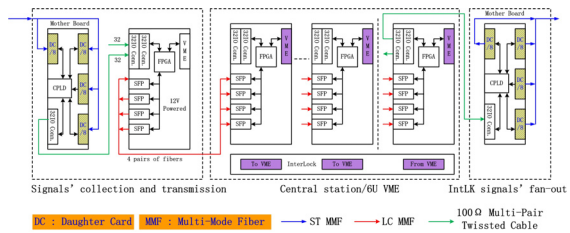


Figure 2: Structure of the FPS.

Signals collection are done locally near the interlock sources, and then transmitted to the central station via high speed serial optical links. Total 16 signals are carried with one pair of fibers. The high speed serial link are implemented with the Rocket I/Os of the XC6SLX100T-3FGG900 from Xilinx and optical transceivers. A dedicated auto-sync logic has been developed for the serial link to work when both fibers of the same link are plugged into the optical transceivers under normal conditions.

Figure 3 shows the picture of the main board. Either VME +5V or external +12V can be used to power the board.

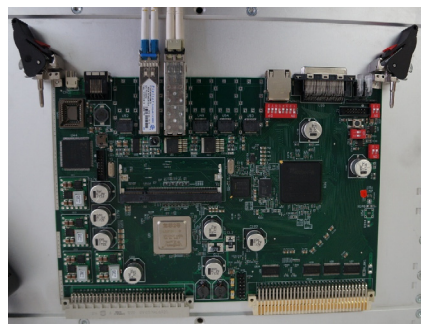


Figure 3: Picture of the main board of FPS.

To improve the overall reliability, a real-time heart-beat function is used for each input signal. That is, heart-beat signals are generated at the collection part, and then checked in the main boards at the central station. A width of 100ns is defined for the heart-beat signals in our case, see Fig. 4.

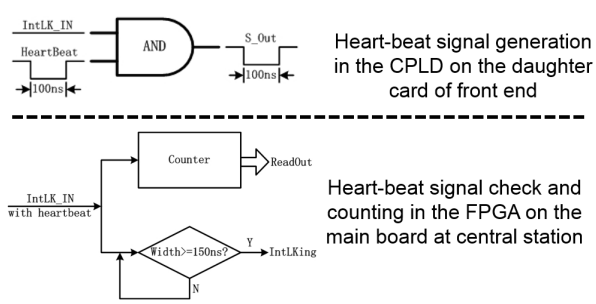


Figure 4: Illustration of the heart-beat function.

*Work supported by National Development and Reform Commission of China, Chinese Academy of Sciences, and Guangdong Province.
#jindp@ihep.ac.cn

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

DEVELOPMENT AND REALISATION OF THE ESS MACHINE PROTECTION CONCEPT

A. Nordt, R. Andersson, T. Korhonen, A. Monera Martinez, M. Zaera-Sanz, European Spallation Source ERIC, Lund, Sweden
A. Apollonio, R. Schmidt, CERN, Geneva, Switzerland
C. Hilbes, ZHAW, Winterthur, Switzerland

Abstract

The European Spallation Source ERIC (ESS) is facing extremely high beam availability requirements and is largely relying on custom-made, specialised, and expensive equipment for its operation. The average proton beam power of 5MW per pulse will be unprecedented and its uncontrolled release can lead to serious damage of the delicate equipment, causing long shutdown periods, inducing high financial losses and, as a main point, interfering drastically with international scientific research programs relying on ESS operation. Implementing a fit-for-purpose machine protection concept is one of the key challenges in order to mitigate these risks. The development and realisation of the measures needed to implement such concept to the correct level in case of a complex facility like the ESS, requires a systematic approach, and will be discussed in this paper.

ESS MACHINE PROTECTION

The European Spallation Source ERIC (ESS), currently under construction, consists of a 600m long high power proton Linac, accelerating proton pulses of 2.86ms length to the energy of 2GeV with a repetition rate of 14Hz. The proton pulses (5MW) then interact with a rotating tungsten target wheel. Neutrons are created due to the spallation process and further guided through 22 different neutron beam lines towards the experimental stations. It is expected to have the first protons on target mid 2019 [1].

As a user facility for neutron science, overall availability of the ESS needs to be defined from a user point of view. Hence, it should be characterized by the average neutron production during a certain time period. ESS availability is interpreted as the average proportion of beam production time during scheduled operation time. In general, the availability characteristics of a system are determined by its reliability, maintainability and inspect-ability. High operational availability is achieved by increasing the mean time between maintenance while avoiding large mean downtimes. A detailed discussion is presented in [2, 3].

In the context of ESS Machine Protection (MP), the term “machine” or Equipment Under Control (EUC) encompasses all elements in the Accelerator, Target Station and Neutron Science system segments - all being necessary for neutron beam production and its further use by the neutron science experiments.

Annika.Nordt@esss.se

ESS MP Design Approach

MP is concerned with operational goals of the ESS, that means, enabling neutron science and investment protection. It is not concerned with safety aspects that are regulated by legal authorities, such as personnel safety or public safety. Nevertheless, because of the high-risk potential associated with damage to the machine, elements related to MP will be implemented in accordance with functional safety standards [4, 5]. To keep the distinction between safety and protection requirements as transparent as possible, adequate definitions for MP have been introduced. The IEC 61511 is defining a Safety Integrity Level (SIL) [5], whilst we refer to the Protection Integrity Level (PIL) for MP instead [6, 7]. Protection Integrity describes the average probability of a protection system satisfactorily performing the required protection functions under all the stated conditions within a stated period of time.

ESS MP Goals and Top Level Requirements

The EUC is exposed to potential damage sources related to proton and neutron beam properties, related radiation, electrical power, vacuum, cooling, RF, etc. The severity of damage is defined with respect to neutron beam quality losses, quality loss duration and resource costs for the recovery of operational capabilities. The goals for MP are defined as follows [6]: MP shall, in that order, prevent and mitigate damage to the machine, be it beam-induced or from any other source, in any operating condition and lifecycle phase, in accordance with beam- and facility-related availability requirements. In addition, MP shall protect the machine from unnecessary beam-induced activation having a potential to cause long-term damage or increase maintenance times.

These two goals can be achieved if MP fulfils the top-level functional and operational requirements being described in the following. MP shall detect all relevant off-nominal states that can lead to damage or unnecessary beam-induced activation and take all the necessary actions for its prevention and mitigation. Protection functions shall be implemented with timing and protection integrity levels in accordance with damage risk reduction requirements. The protection functions shall be implemented such that the probability of spurious trips is reduced. All information about detected off-nominal states, performed prevention and mitigation actions shall be recorded, allowing for an a-posteriori reconstruction and analysis. Operation during all foreseen lifecycle phases of the machine, for example assembly and

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

REUSABLE PATIENT SAFETY SYSTEM FRAMEWORK FOR THE PROTON THERAPY CENTRE AT PSI

P. Fernandez Carmona, M. Eichin, M. Grossmann, E. Johansen, A. Mayor, H.A. Regele
PSI, Villigen, Switzerland

Abstract

A new gantry for cancer treatment is being installed at the Proton Therapy Centre in the Paul Scherrer Institut (PSI), where already two gantries and a fix line operate. A protection system is required to ensure the safety of patients, requiring stricter redundancy, verification and quality assurance measures than other accelerators. It supervises the Therapy System, sensors, monitors and operator interface and can actuate magnets and beam blockers. We built a reusable framework to increase the maintainability of the system using the commercial IFC1210 VME controller, developed for other PSI facilities. It features a FPGA implementing all the safety logic and two processors, one dedicated to debugging and the other to integrating in the facility's EPICS environment. The framework permitted us to reduce the design and test time by an estimated 40% thanks to a modular approach. It will also allow a future renovation of other areas with minimum effort. Additionally it provides built-in diagnostics such as time measurement statistics, interlock analysis and internal visibility. The automation of several tasks reduces the burden of QA in an environment with tight time constraints.

INTRODUCTION

The Paul Scherrer Institut (PSI) was a pioneer in the field of proton therapy for cancer treatment by being the first centre to implement spot scanning for dose delivery back in 1996. Nowadays there exist many centres using such technology and several vendors offering commercial products. In order to increase the number of patients being treated it was the decided to buy a scanning gantry from Varian Medical Systems [1], while keeping research and development in the existing in-house engineered areas. The current facilities of the CPT consist of a fixed beam line for eye cancer treatment, operating clinically since 2010, and Gantries 1 and 2 operating since 1996 and 2013 respectively [2] - [3]. The beam is provided by a dedicated 250 MeV cyclotron from the company Varian Medical Systems.

There are several systems required to allow for a safe and accurate delivery of the prescribed dose to the patient. The most relevant ones are the Patient Safety System (PaSS), to prevent accidents and the Therapy Control System (TCS), to deliver dose and to verify the correctness of the delivery. Other systems working independently but interconnected are the Beam Tuning Verification System (BTVS), the Machine Control System (MCS), the Run Permit System (RPS) and the Main Patient Safety Switch and Controller (MPSSC). Finally there is a number of Dose and Beam Position Monitors and a set of final elements such as a kicker magnet and beam blockers.

In order to integrate the commercial gantry in the existing facility it was necessary to develop two adapters: The TCS adapter interfaces the vendor specific control system commands (such as setting energy and beam current values) to the appropriate facility resources. The PaSS adapter autonomously takes care of preventing accidents, and also actuates some final elements on request of the gantry. Figure 1 shows an overview of all the systems to which the PaSS is connected.

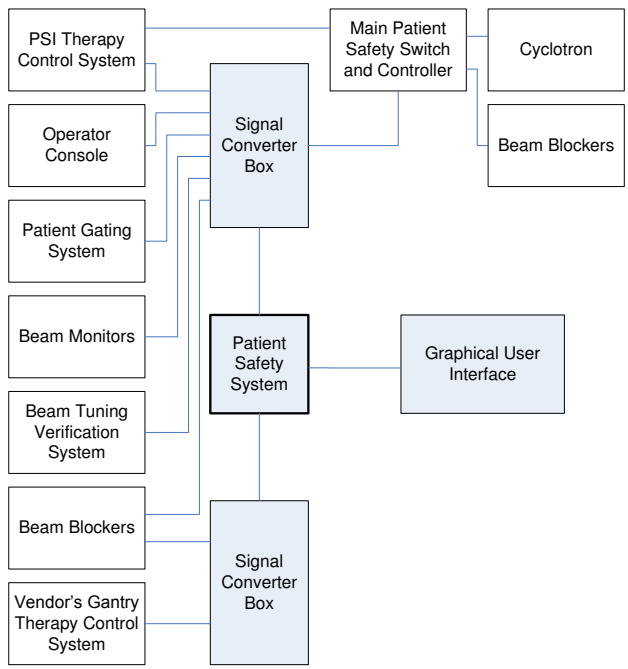


Figure 1: System overview of the integration of the new gantry's Patient Safety System.

This document describes in detail the new Patient Safety System developed for the installation of the new gantry. Also the experience of conceiving and designing it as a reusable framework, comparing it to previous designs is presented. Final tests for regulatory approval and clinical commissioning will take place during 2016 and patient treatment will start later in the year.

PATIENT SAFETY SYSTEM CONCEPT

The main goal of the patient irradiation for CPT is that each dose spot is delivered at the correct position and with the correct dose quantity. According to the International Commission on Radiation Units and Measurements (ICRU) [4], the aim in intensity modulated radiation therapy is to achieve an accuracy of 5% of the

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

NSLS-II ACTIVE INTERLOCK SYSTEM FOR FAST MACHINE PROTECTION*

K. Ha#, S. Seletskiy, Y. Tian, Y. Hu, W. Cheng, O. Singh, D, Padrazo, A. Castiblanco, J. Mead, R. Smith, Y. Tang, W. Guimei, P. Ilinski, J. De Long, L. Dalesio BNL, Upton, NY 11973, USA
G. Shen, FRIB, East Lansing MI 48824 USA

Abstract

At National Synchrotron Light Source-II (NSLS-II), a field-programmable gate array (FPGA) based global active interlock system (AIS) has been commissioned and used for beam operations. The main propose of AIS is to protect insertion devices (ID) and vacuum chambers from the thermal damage of high density synchrotron radiation power. This report describes the status of AIS hardware, software architectures and operation experience.

INTRODUCTION

NSLS-II synchrotron radiation source produces very wide spectrum ranges from IR very hard x-ray. Insertion devices (IDs) and front end devices commissioning started in November 2014 and user operation started in February 2015. The 6 NSLS-II project beamlines consist of EPU, IVU, VUV IDs. The AIS is a key machine protection system that protects vacuum chambers and front end devices from synchrotron radiation power. NSLS-II storage ring (SR) has 792 m circumference and consists of 30 double-bend achromat cells. 180 SR RF beam position monitors (BPMs) are installed around the ring, and 23 ID BPMs are installed in the straight sections. They are all fully functional. The AIS was installed in January 2014 and was successfully commissioned. It has been used in daily operation since November 2014. Recent commissioning status and results of AIS were well described in [1]. NSLS-II AIS has a state-of-the-art architecture that include 30 cell controllers (CCs) and one AIS controller based on a Xilinx Virtex-6 FPGA. The requirement for AIS is to dump the stored beam using RF interrupt within 1ms. It prevents increase of chamber temperature when orbit is mis-steered. We achieved 22 µs global BPM data transfer latency around the ring through 1 to 31 CC nodes. Each of cell controllers contains 5 Gbps transceivers running home-made serial device interface (SDI) protocol. All calculation logic is implemented using FPGA. During beam commissioning, functional test was fully completed with stored beam (1 mA ~ 300 mA) with ID gap open/close, and with ID and bending magnet (BM) photon shutters open/close. The SDI link 10 kHz packet communication and angle offset calculation results were very stable and fast enough for commissioning and user operation. In addition, the measured actual beam dump time is about 1 ms and system hardware latency time is about 200 µs.

*Work supported by DOE contract No: DE-AC02-98CD10886.
#kha@bnl.gov

SYSTEM OVERVIEW

Our novel architecture of AIS is shown in Figure 1. AIS employed a 1 AIS controller, 30 CCs, 203 BPMs, 1 DCCT, 8 front end PLCs. In each cell, local SDI link connected 6 RF BPMs and additional 3 ID BPMs. AIS controller is installed in cell 23 rack group D. Each cell has its own PLC for the integration of equipment protection system (EPS) logic, for ID control and photon shutter controls. This PLC provides ID gap status, BM photon shutter status and ID photon shutter status to the local CC. Each CC communicates with its neighbor cell via 40-m multi-mode fibre optics cables. A ring topology is established by the 30 CCs with SDI protocol. The BPM position data, PLC data, and DCCT status data are delivered to the AIS controller at 10 kHz rate.

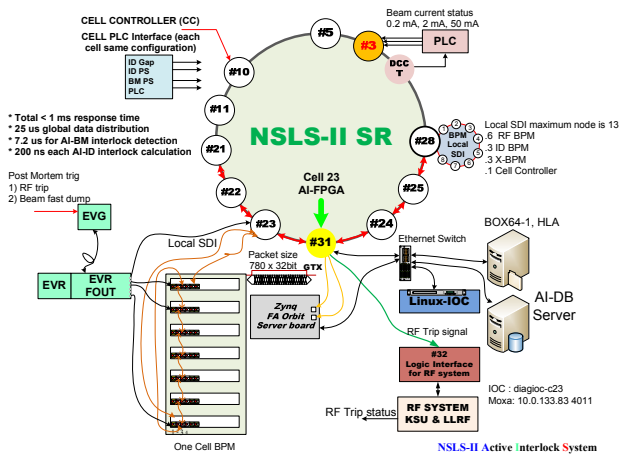


Figure 1: AI system installed layout. (Cell23: AIS and IOC installed, Cell3: DCCT controller installed, Project beamline (Cell 3, 5, 10, 11, 23, 28, 8, 18), and ABIX beamline (16, 17) IDs installed)

The AIS specification is developed following the NSLS-II synchrotron radiation protection requirements [2]:

1. Support multiple IDs and BMs.
2. The AIS controller must cover up to 64 IDs.
3. 10 kHz system monitoring and calculations.
4. Easily extensible for additional IDs which will be installed in the future beamlines.
5. Total system response time is less than 200 µs (BPM to RF trip).
6. Redundant communication.
7. Maximum fast beam dump latency time is less than 1 ms.
8. Slow beam dump time is less than 10 ms (LLRF power ramp down).

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

MACHINE PROTECTION SYSTEM FOR THE KOMAC 100-MeV PROTON LINAC

Young-Gi Song, Kyung-Tae Seol, Han-Sung Kim, Dae-il Kim, Sang-Pil Yun, Hyeok-Jung Kwon, Yong-Sub Cho, Korea Multi-purpose Accelerator Complex, Korea Atomic Energy Research Institute, Korea

Abstract

A Machine Protection System (MPS) is one of the important systems for the 100-MeV proton linear accelerator of the Korea Multi-purpose Accelerator Complex (KOMAC). The MPS is required to protect the very sensitive and essential equipment during machine operation. The purpose of the MPS is to shut off the beam when the Radio-Frequency (RF) and ion source are unstable or a beam loss monitor detects high activation. The MPS includes a variety of sensors, such as beam loss, RF and high voltage converter modulator faults, fast closing valves for vacuum window leaks at the beam lines and so on. The MPS consists of a hardwired protection for fast interlocks and a software protection for slow interlock. The hardware-based MPS has been fabricated, and the requirement has been satisfied with the results within 3 μs. The Experimental Physics and Industrial Control System (EPICS) control system has been also designed to monitor and control the MPS using a Programmable Logic Controller (PLC). This paper describes the design and implementation of the MPS for the 100-MeV proton linear accelerator of the Korea Multi-purpose Accelerator Complex (KOMAC).

KOMAC FACILITY OVERVIEW

The Korea multi-purpose accelerator complex (KOMAC) 100-MeV proton linac has been developed and has been installed at the Gyeong-ju site. Figure 1 shows a schematic layout of the KOMAC 100-MeV proton linac and beam lines [1].

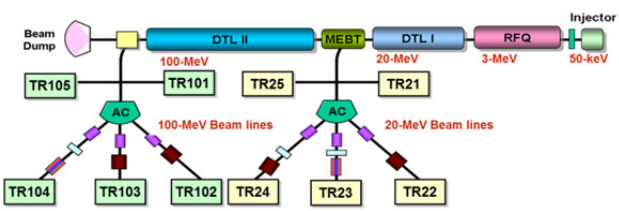


Figure 1: KOMAC proton linac consists of an ion source, a Low-Energy Beam Transport (LEBT), a Radio-Frequency Quadrupole (RFQ), a Draft Tube Linac (DTL), a Medium-Energy Beam Transport (MEBT), beam-lines, and a Target Room (TR) for 20-MeV and 100-MeV beam.

The KOMAC consists of low-energy components, including a 50-keV ion source, a low-energy beam transport (LEBT), a 3-MeV Radio-Frequency Quadrupole (RFQ), and a 20-MeV Drift Tube Linac (DTL), as well as high-energy components, including seven DTL tanks for the 100-MeV proton beam. The KOMAC includes 10 beam lines, 5 for 20-MeV beams and 5 for 100-MeV beams. The peak beam current and maximum beam duty of the 20-MeV linac are 20 mA and 24%, respectively. The peak beam current and maximum beam duty of the 100-MeV linac are 20 mA and 8%, respectively. There are 4 high-voltage convertor modulators. Each modulator drives 2 or 3 klystrons. The peak output power is 5.8 MW, and the average power is 520 kW with a duty of 9%. The pulse width and the repetition rate are 1.5 ms and 60 Hz, respectively.

MPS OVERVIEW

The radiation from the beam loss and faults of the linac components can cause substantial damage to the devices. The KOMAC active protection system needs to minimize the beam loss radiation and ensure the safe operation of the machine [2]. The purpose of a machine protection system (MPS) is to turn off the beam and the sub-systems when an interlock occurs. The MPS consists of two parts: a hardwired-based fast interlock system using an analog circuit and a software-based slow interlock system with a Programmable Logic Controller (PLC). The interlock systems are designed as main unit for machine protection and are fabricated with analog circuits. The PLC includes software logic that checks the result of the machine protection system on an interlock signal and also turns off components related to the interlock. Figure 2 describes a flow chart of machine protection using hardwared and software based interlock system.

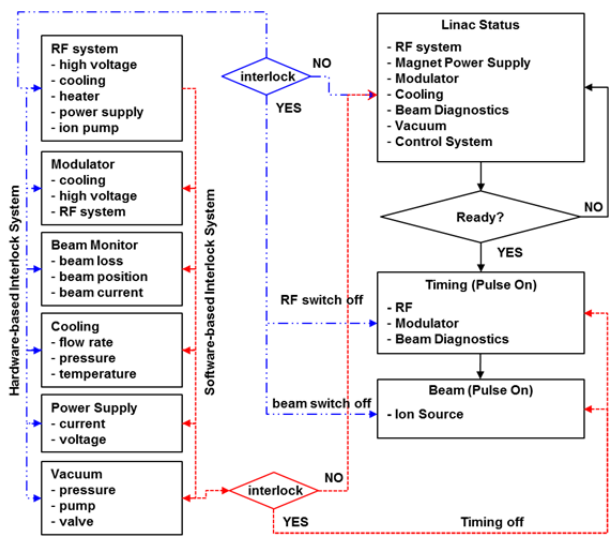


Figure 2: Block diagram of interlock system for machine protection with hardware and software based interlock system.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

SAFETY INTEGRITY LEVEL (SIL) VERIFICATION FOR SLAC RADIATION SAFETY SYSTEM

F. Tao, E. Carrone, J. Murphy, K. Turner
SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

Abstract

Safety Integrity Level (SIL) is a key concept in functional safety standards. SIL is a performance measure on how reliable a safety system is in performing a particular safety function. To comply with standards, during the system design stage, SIL verification must be performed to demonstrate that the SIL achieved meets/exceeds the SIL that was assigned during risk assessment. Unlike industrial applications, where safety systems are usually composed of certified devices or devices with long usage history, safety systems in large physics laboratories are less standardized and more complex in terms of system architecture and devices used. In addition, custom designed electronics, with limited reliability information, are often employed. Verifying SIL for these systems requires in-depth knowledge of reliability evaluation. In this paper, it is demonstrated how to determine SIL using SLAC radiation safety systems (Personnel Protection System (PPS) and Beam Containment System (BCS)) as examples. PPS utilizes commercial safety rated devices, while BCS contains customized electronics. Choice of standards, methods of evaluation, reliability data gathering process (both from industry and from hardware development) are also discussed.

INTRODUCTION

At SLAC National Accelerator Laboratory, there are three protection systems deployed to mitigate radiation risks: Machine Protection System (MPS), Beam Containment System (BCS) and Personnel Protection System (PPS). While MPS is focused on protecting equipment from getting damaged, the other two systems are critical to protect personnel and the environment from getting a radiation dose. In comparison with the MPS, PPS and BCS have very rigorous configuration control policies in place and follow all due diligence in engineering practices. For these reasons, these two radiation safety systems meet all definitions of safety-critical control systems.

Functional safety standards started in 1990's and now there are several critical standards have been developed and adopted worldwide. Such as IEC61508 [1] and corresponding sector specific standards IEC61511 [2] (process), IEC 62061 [3] (machinery). ISO also published a machinery safety standard ISO13849 and its relationship with IEC62061 is described in a IEC/ISO joint technical report ISO/TR23849 [5]. While industries have started following functional safety standards in implementing safety-critical control systems two decades ago, the natural question that arises is, are the same standards applicable to accelerator safety systems, and if the answer is yes, how are the standards applied so that we can learn from industries' long time experience with the engineering design and operation of safety systems.

As found in any functional safety standard, two key concepts are safety lifecycle and safety integrity level. The former describes a series of activities in system engineering and operation to make sure all risks are identified and properly mitigated. The latter is a measure of the reliability performance each safety function within the safety system, such that the design be precisely performance based and is less conservative than the traditional description-based approach.

In all functional safety standards, after the conceptual design stage, the SIL of the safety function must be verified to make sure it meets or exceeds the SIL assigned in the Safety Requirements Specification (SRS). In the case the safety system/function SIL level is not met, the safety system/function must be re-designed. For example, the following figure is from IEC 61511.

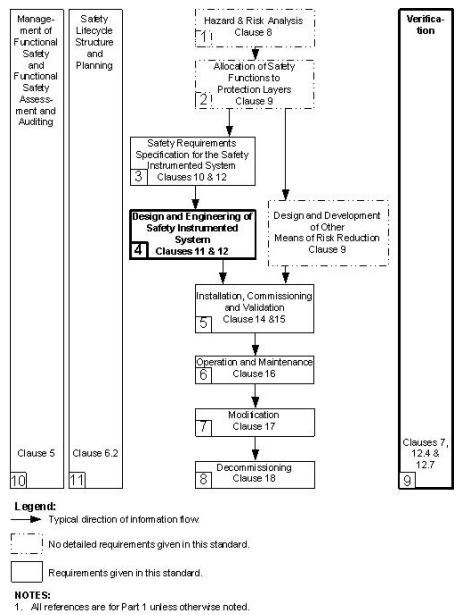


Figure 1: Safety System Lifecycle Phases.

Though PPS and BCS are both classified as safety-critical systems with the purpose to mitigate radiation risks, different philosophies and technologies are used. The majority of the PPS devices are commercial off-the-shelf (COTS), their reliability data can be obtained and the system level reliability block diagram (RBD) can be

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

THE LMJ TARGET DIAGNOSTICS CONTROL SYSTEM ARCHITECTURE

S. Perez CEA/DIF, Bruyères le Châtel, 91297, Arpajon, France
T. Caillaud, CEA/CESTA, Le Barp, 33114, France

Abstract

The French Laser Megajoule (LMJ) is, behind the US NIF, the second largest inertial fusion facility in the World. More than 30 diagnostics will be installed and driven in a huge and complex integrated computer control system. The aim of this paper is to describe an architecture based on the TANGO open source software for the very low level control system, Python language for the development of drivers and the French commercial PANORAMA© software as the main high level SCADA.

This choice leads to guaranty the evolution of the middleware software architecture of this facility supposed to be operated during dozen of years with the capability of using many instruments including sustainability.

INTRODUCTION

Since it definitively abandoned nuclear testing, France has relied on the Simulation Program to guarantee the operational performance and safety of its nuclear deterrent weapons throughout their lifetime.

Successful simulation requires both:

- Qualified computer codes that integrate laboratory-validated physics models to simulate weapon functioning;
- Teams of qualified physicists to use these codes.

In this respect, the Megajoule Laser (LMJ [1]) plays a vital role, as it is used to validate the numerical codes and certify the skills of French physicists.

On October 23, 2014, French Prime Minister Manuel Valls declared the facility operational after starting up the first experiment.

Target diagnostics are a key for numerous physical data acquisition. CEA will develop dozen of these equipments during next twenty years. Each target diagnostic will be dedicated to one or several kind of measurements like X-ray, visible, UV or particles like neutron…

During the life cycle of the LMJ Facility, CEA needs to implement new kind of equipments compliant with a stabilized control command system.

This paper describes the command control architecture used for target diagnostics and the reasons why we insure sustainability for such a huge modular installation.

We will describe the 2 layers of the Target Diagnostics Control System (TDCS) and particularly the use of TANGO for Layer 0, which guaranty modularity and life time expectancy, and the French SCADA PANORAMA E² [2] for Layer 1 dedicated for every LMJ command control subsystem.

Maintenance and qualification tools will also be described. The use of Open Source Software like TANGO, Python and QT will allow the capability for diverse contractors to insure all future developments.

USING A TARGET DIAGNOSTIC

In 2014-2015, three different target diagnostics have been installed: two X-ray imaging systems (different ranges of energy and waveform) and a complex diagnostic used for Hohlraum temperature measurements including an absolutely calibrated broadband X-ray Spectrometer, a Gating Spectrometer and a time resolved Imaging System of the emitting area.

This paper will focus on a « simple » X-ray imaging system. This diagnostic can actually be compared to a giant microscope. Made of 4 parts (alignment beams, a telescopic motorized arm, filters and a framing camera), the command control configures the equipment, then focus it using alignment and the arm, in order to acquire data from an optical camera (Fig. 1).



Figure 1: an X-Ray Target Diagnostic.

Each part of the diagnostic has to be driven by the command control and, as these functions should be reused in future target diagnostics, a modular command control architecture must be designed.

THE LMJ COMMAND CONTROL ARCHITECTURE

The LMJ command control architecture is driven by the 4 classical component layers, as shown in Figure 2 :

- Layers 2 and 3 are devoted to the common control system (administration, main supervisory, prediction and tuning system [4] [5], sequences [6]…)
- Layer 1 is set for the main subsystems command controls (target and laser diagnostics, synchronization [7]…) and interfaces between them.
- Layer 0 is the main layer for equipment communications. It includes drivers, communication protocols as well as maintenance and qualification tools.

There is one Layer 1 for the all TDCS and as many as necessary drivers included in Layer 0 for each defined

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

EXTREME LIGHT INFRASTRUCTURE, BEAMLINES – CONTROL SYSTEM ARCHITECTURE FOR THE L1 LASER*

J. Naylon#, T. Mazanec, Institute of Physics, Prague, Czech Republic, A. Greer, C. Mayer, Observatory Sciences Ltd., Cambridge, UK, M. Horáček, B. Himmel, M. Drouin, K. Kasl, J. Horáček, P. Škoda, P. Bakule and B. Rus, Institute of Physics, Prague, Czech Republic

Abstract

The ELI-Beamlines facility aims to provide a selection of high-energy and high repetition-rate TW-PW femtosecond lasers driving high intensity XUV/X-ray and accelerated particle secondary sources for applications in materials, medical, nuclear and high-field physics sectors. The highest repetition rate laser in the facility will be the L1 laser, producing 1 kHz, 20 fs laser pulses of 200 mJ energy. This laser is based entirely on picosecond chirped-pulse parametric amplification and solid-state pump lasers. The high repetition rate combined with kW pump powers and advanced technologies calls for a highly automated, reliable and flexible control system. Current progress on the L1 control system is discussed, focussing on the architecture, software and hardware choices. Special attention is given to the LabVIEW-EPICS framework that was developed for the ELI Beamlines lasers. This framework offers comprehensive and scalable EPICS integration while allowing the full range of LabVIEW real-time and FPGA embedded targets to be leveraged in order to provide adaptable, high-performance control and rapid development.

INTRODUCTION

The four laser sources (L1-L4) are currently under development and will be installed in the facility in time for a 'first light' demonstration in 2016-17. The L1 laser's high repetition rate, 5 TW peak power and excellent beam quality is aimed at experiments based on high-harmonic generation (XUV), X-ray and keV betatron radiation in the molecular, biomedical and materials sciences [1].

The resulting L1 control system architecture and design is currently in a fairly final state with most hardware and software components operational and tested as part of an integrated solution. The aim of this paper is to summarise the chosen architecture and highlight some of its advantages for similar control system projects.

ARCHITECTURE DEVELOPMENT

At the start of control system development, the type and quantity of I/O was estimated by reference to similar laser systems. Requirements were also collected regularly from the laser scientists. The estimate of the scale of the I/O requirements was the basis of the architecture:

Table 1: L1 Control System Scale

Device	Parameters	Approx. qty
Cameras	20 Hz, 2 M pixel	60
High speed cameras	1 kHz 0.5 M pixel	6
Stepper motors	Simple, open loop	100
Piezo actuators	Small, open loop	20
Piezo steppers	Mostly open loop	20
Laser relative energy	*e.g.,* photodiode, kHz	40
Digital I/O channels	Mostly 24 V logic	30
General analogue I/O	≤kHz, 16 or 24-bit	30
Temperature sensors	RTDs	20
Flow sensors	Cooling water, pulse	20
Complex instruments	Serial, USB, Ethernet	6
Simple devices	Serial, Modbus, *etc.*	15

Over the course of development other project-specific factors have also been significant in determining the final choice of implementation:

- The high laser repetition rate and expense of key laser optics necessitates a real-time control system with sufficient intelligence to provide machine protection;
- Some laser instruments are highly specific and available from few vendors. Generally the only integration for these is via LabVIEW;
- The laser is based on new and cutting-edge technology, so control system requirements change frequently;
- The budget, time and effort required for laser control systems are often underestimated;
- Human resources, particularly software developers, are limited due in part to competition from the rapidly growing IT sector locally;
- Being a new facility there is little established experience. Any solution must be easy and quick for new staff to learn and work with;
- Strict tendering rules and laws restrict procurement; therefore reliance on a few key suppliers for most of the control system components justifies the time and expense of preparing a framework contract.

EPICS was chosen as the integration framework as it is widely used, stable and requires minimal programming. Use of LabVIEW was essential and is fast, flexible and easy to learn. This was chosen to be the sole software development language for reasons of maintainability and to streamline training. The resulting structure of the control system (Fig. 1) is based on a simple 3-layer scheme appropriate to the scale and suitable for an industrial/machine-control system.

* Work supported by the European Regional Development Fund and the European Social Fund under Operational Programs ECOP and RDIOP. The project is hosted by Fyzikální ústav AVČR, v.v.i.
#jack.naylon@eli-beams.eu

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

REMUS: THE NEW CERN RADIATION AND ENVIRONMENT MONITORING UNIFIED SUPERVISION

Adrien Ledeul, Gustavo Segura Millan,
Riku-Pekka Ilari Silvola, Bartlomiej Styczen, Daniel Vasques Ribeira,
CERN, Geneva, Switzerland

Abstract

The CERN Health, Safety and Environment Unit is mandated to provide a Radiation and Environment Monitoring SCADA system for all CERN accelerators, experiments as well as the environment. In order to face the increasing demand of radiation protection and continuously assess both the conventional and the radiological impact on the environment, CERN is developing and progressively deploying its new supervisory system, called REMUS - Radiation and Environment Monitoring Unified Supervision.

This new WinCC OA based system aims for an optimum flexibility and scalability, based on the experience acquired during the development and operation of the previous CERN radiation and environment supervisory systems (RAMSES and ARCON). REMUS will interface with more than 70 device types, providing about 3,000 measurement channels (approximately 500,000 tags) by end 2016.

This paper describes the architecture of the system, as well as the innovative design that was adopted in order to face the challenges of heterogeneous equipment interfacing, diversity of end users and non-stop operation.

INTRODUCTION

Radiation Protection and Environment Monitoring of CERN facilities and experiments are essential for three reasons:

- Safety of the workplace and of the environment.
- Reporting to authorities the nature and the quantities of emitted ionizing radiation.
- Reporting to authorities in case of pollution of the environment.

In order to fulfil those missions, CERN has set up diverse monitoring equipment across CERN area and its immediate surroundings.

Since 2005, RAMSES (RAdiation Monitoring System for the Environment and Safety) [1] has been in charge of the data acquisition, control/command and supervision of 50 device types, allowing the supervision of 1,500 channels. The former radiation monitoring system, ARCON (ARea CONtroller), is also still in charge of the majority of the monitoring in PS (Proton Synchrotron) accelerator complex. Additionally, proprietary supervisors such as Berthold MEVIS handle the data acquisition and supervision of devices that have not been integrated into RAMSES Supervision. Numerous devices are also operated without any remote supervision.

In 2012, the development of a new SCADA system, REMUS (Radiation and Environment Monitoring Unified Supervision) began. This system aims to:

- Cover a larger set of equipment than RAMSES, in order to unify the supervision and integrate stand-alone devices, most of them being COTS (Commercial Off-The-Shelf) products.
- Provide a reliable, scalable and cost-effective system.
- Use common CERN software, WinCC OA (WinCC Open Architecture) [2], formerly known as PVSS, and JCOP framework (Joint COntrols Project) [3], for the development and support of the new SCADA.
- Reduce the delay and the cost of adding new devices to the supervision.
- Provide light and fast clients, adapted to the needs of diverse end-users.
- Reduce the overall maintenance and user support effort necessary to maintain the system in operation.

REMUS founded its evaluation of scalability and performance on the study performed by CERN engineering department [4]. The probed technologies are commonly employed in large SCADA systems supervising LHC (Large Hadron Collider) experiments.

For its development, REMUS applies the same International Standard IEC 61508 [5] that was set up during the development of RAMSES.

REQUIREMENTS AND SCOPE

The REMUS project aims at developing a universal software for supervision, control and data acquisition for the entire suite of monitoring stations covering all radiological and environmental parameters that can potentially be affected by the operation of the facilities of CERN.

The functions of the system include:

- Logging of all measured values and system events coming from the instrumentation.
- Providing real-time measured values, alarms and operational states of the devices through customizable user interfaces composed of synoptic and widgets.
- Publishing alarms and faults of the devices to a remotely accessible user interface and to the CERN central alarm system, LASER (Lhc Alarm SERvice) [6].
- Publishing a selected set of measurements to other systems via the CERN Data Exchange system, DIP (Data Interchange Protocol), built on the top of DIM (Distributed Information Management) [7].

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

THE VIRTUAL EUROPEAN XFEL ACCELERATOR

R. Kammering, W. Decking, L. Froehlich, O. Hensler, T. Limberg, S. Meykopff, K. Rehlich, V. Rybnikov, J. Wilgen, T. Wilksen, DESY, Hamburg, Germany

Abstract

The ambitious commissioning plans for the European XFEL require that many of the high-level controls are ready from the beginning. The idea arose to create a virtual environment to carry out such developments and tests in advance, to test interfaces, software in general and the visualisation of the variety of components. Based on the experiences and on the systems that are already in operation at the FLASH facility for several years, such a virtual environment is being created. The system can already simulate most of the key components of the upcoming accelerator. Core of the system is an event synchronized data acquisition system (DAQ). The interfaces of the DAQ system towards the device level, as well as to the high-level side is utilising the same software stack as the production system does. Thus, the software can be developed and used interchangeably between the virtual and the real machine. This allows to test concepts, interfaces and identify problems and errors at an early stage. In this paper the opportunities arising from the operation of such a virtual machine will be presented. The limits in terms of the resulting complexity and physical relationships will also be shown.

THE IDEA

Lessons learned from the fast successful start up of the Linac Coherent Light Source (LCLS) at the SLAC National Accelerator Laboratory showed that one needs to have not only all hardware related software ready and checked from the first hours on, but also all foreseen high level software.

Since most high level software is not directly acting on the hardware layer itself but needs a vast amount of infrastructure to be working properly, testing high level software can only be done with this infrastructure up and running. However, having the high level software ready and tested prior to having the first beam is in contradiction. To still be able to accomplish this task, the idea of having a test environment for this, came up within our group which is in charge of creating high level controls and applications. For a general overview of the concepts and architecture of the foreseen high level software for the European XFEL see [1].

FROM FLASH TO XFEL

The FLASH facility can be seen as the *small brother* of the European XFEL. On the one hand due to the use of the same hardware technologies, but also in the sense of layout, beam physics and creation of the FEL pulses.

Similarities

FLASH has always been used for prototyping and testing hardware to be later be used at the European XFEL. To name some examples of the major hardware systems, there are the:

- µTCA crate standard
- Timing System
- Machine Protection System

Thus FLASH serves already since several years as the test bed for the interplay between the new hardware systems and the software layers for processing and providing this data to the operators and experts. But the European XFEL has, roughly speaking, ten times the extend compared to the FLASH facility. Table 1 shows a comparison of some hardware and characteristics for both machines.

Table 1: FLASH vs. XFEL

System	FLASH	XFEL
Crates	~ 30 VME, µTCA crates	~ 200 µTCA crates
BPMs	~ 40	~ 460
Data rate to DAQ	< 100 Mbytes/s	>> 100 Mbytes/s
Length	~ 300 m	~ 3000 m

The shear number of devices to be installed in the European XFEL and the thereby resulting data rates require a strong data reduction and pre-processing already on the lower hard- and software layers. But beside the need for data concentration at an early stage, the operation of such a large scale facility as the European XFEL require precise synchronization of many different hardware devices.

Further did lessons learned from the operation of the FLASH facility show that one needs to work more in terms of physical meaningful values instead of reading *raw* data directly from the hardware devices.

Even though many ideas and concepts for high level applications can be lend from the FLASH facilty, one needs to take the increased data rate into account. A central concept to overcome this, is to use the, at the FLASH facility already since years used, data acquisition system.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

INTEGRATING CONTROL APPLICATIONS INTO DIFFERENT CONTROL SYSTEMS*

M. Killenberg†, M. Hierholzer, C. Schmidt, DESY, Hamburg, Germany
S. Marsching, aquenos GmbH, Baden-Baden, Germany
J. Wychowaniak, Łódź University of Technology, Łódź, Poland

Abstract

Porting complex device servers from one control system to another is often a major effort due to the strong code coupling of the business logic to control system data structures. Together with its partners from the Helmholtz Association and from industry, DESY is developing a control system adapter. It allows writing applications in a control system independent way, while still being able to update the process variables and react on control system triggers. We report on the status of the project and the experience we gained trying to write portable device servers.

INTRODUCTION

With embedded systems becoming more and more powerful, the algorithms in the devices which are accessed via control systems are becoming more and more advanced. Especially on MicroTCA [1] systems the hardware usually features a powerful multi-core CPU with several GB of RAM. The MicroTCA.4 [2] extension brings trigger and clock lines, as well as large rear transition modules which can be used for demanding analogue control applications.

Many particle accelerators which are currently being build are using or will use MicroTCA.4 for control of the radio frequency (RF) in the accelerator, for instance FLASH [3] and the European XFEL [4] hosted at DESY, Hamburg, ELBE [5] at Helmholtz-Zentrum Dresden-Rossendorf, or FLUTE [6] at KIT, Karlsruhe. The complex RF control applications shall be reused across the different accelerators, while all the facilities are using different control systems. It turned out that porting the software to a different control system is a major effort because the code is strongly coupled to the original control system. This lead to the idea to have an adapter layer between the device library, which implements the algorithms, and the control system, which provides the communication protocols and integration into the facility's control infrastructure. This adapter shall be part of the MicroTCA.4 User Tool Kit (MTCA4U) [7, 8], a collection of libraries to facilitate the implementation of control applications.

REQUIREMENTS

The main task of the adapter is to allow application code to access process variables which are communicated to the outside world in a control system independent way. For this, the adapter has to use the functionality which is provided by the control system, like communication protocols or the addressing scheme.

The part of code which is device *and* control system dependent has to be minimal, zero if possible. This type of code is causing the huge workload when porting and maintaining applications for multiple control systems.

Abstraction is easy to achieve if data is simply copied back and forth between two domains, but this comes with a performance penalty. So an additional requirement to the adapter is to avoid unnecessary copying, especially of large data structures like arrays. In addition, the adapter should be thread safe and usable in real time application.

As a starting point and to check if the abstraction is working, the adapter is tested with two control systems: DOOCS [9] and EPICS [10]. DOOCS (used at DESY for FLASH and the European XFEL), has an object-oriented data model written in C++. EPICS 3, one the most widely used control systems for particle accelerators and used at FLUTE, has a channel-based C API. We intentionally used two conceptually different control systems, hoping that the abstraction needed to work with these two should also allow other control systems to be used without too much modifications in the design.

DESIGN CONCEPT

The first implementation of the adapter focuses on process variables. It provides data structures for scalars (8, 16 and 32 bit signed and unsigned integers, single and double precision floating point), strings, and arrays of the numerical data types. Each process variable is identified by a unique name which describes its function inside the device code ("TEMPERATURE" for instance for a device with a temperature sensor). The name does not contain information where the device is installed and in which context it is used. This part depends on the control system and the facility, and is added in the control system specific part of the adapter, not in the device part.

The original idea to avoid copying, especially for large arrays, was to have a single instance of the data. This would be stored in a control system dependent type, an instance which always has to be there to work with the particular control system. The adapter would provide a wrapper, which would be used inside the business logic. But it turned out that this approach is not viable. DOOCS and EPICS have different locking schemes for their variables, and only the control system side could know when it is safe to access

* This work is supported by the Helmholtz Validation Fund HVF-0016 "MTCA.4 for Industry".
† martin.killenberg@desy.de

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

THE TANGO CONTROLS COLLABORATION IN 2015

Andrew Götz, Jean-Michel Chaize, Emmanuel Taurel, Pascal Verdier,
Jean-Luc Pons, Tiago Coutinho, Faranguiss Poncet, Reynald Bourtemburg, ESRF, Grenoble, France
Gwenaelle Abeille, SOLEIL, Gif-sur-Yvette, France, Sandor Brockhauser,
Lajos Jeno Fulop, ELI-ALPS, Szeged, Hungary, Alejandro Vázquez-Otero, ELI-BEAMS, Prague, Czech Republic
Mihail Octavian Cernaianu, IFIN-HH, ELI-NP, Bucharest - Magurele, Romania
Cristina Knapic, Ricardo Smareglia, INAF, Trieste, Italy
Igor Khokhriakov, Helmholtz-Zentrum Geesthacht, Geesthacht, Germany

Abstract

This paper presents the latest news from the TANGO collaboration. TANGO is being used in new domains. The three ELI pillars - ELI-Beamlines, ELI-ALPS and ELI-NP in Czech Republic, Hungary and Romania respectively have selected TANGO for many of their control systems. TANGO will be extended with new features required by the laser community. These features will include nanosecond time-stamping. The latest major release of TANGO V9 includes the following features - data pipes, enumerated types, dynamic commands and forwarded attributes. A prototype REST API (mtango) has been developed. The collaboration has been extended to include the new members and to provide a sustainable source of resources through a collaboration contract. A new website (http://www.tango-controls.org/) has been designed which has improved the communication within the community. Finally the updated roadmap is summarised.

THE COLLABORATION

New Sites

The TANGO Controls collaboration has continued to grow. New sites since the last ICALEPCS in 2013 are ELI-BEAMS, ELI-ALPS, ELI-NP, VIRGO EGO and SKA to mention a few. In 2015 two new TANGO-based synchrotrons have been commissioned – Solaris (Poland) and MAX IV (Sweden).

Student Trainings

The number of students attending classes at high schools, universities and training courses has increased. TANGO courses are now available as part of a Bachelor's degree in Grenoble (France), and as part of a Master's degree offered by the University of Szeged (Hungary). TANGO is used for training courses at the EPFL in Lausanne (Switzerland). More students learning TANGO means it is possible to find trainees for TANGO related projects. A group of French students submitted a TANGO project running on a Zynq platform to the Open Hardware 2015 Xilinx student competition. The IAEA has sponsored 3 training courses on TANGO as part of a further education programme in Northern Africa.

TANGO V9

The TANGO approach to development is to continue improving the code base by making regular major releases. The latest major release is version 9. It has been in beta test at the ESRF since October 2014. The official source code release 9.1.0 was made in October 2015.

Why is it a major release? Firstly because of the addition of a new communication channel – pipes. Secondly because there is a new device interface (Device_5Impl). Thirdly because of new implementations in Java for device servers and in Python with a more Pythonic interface. As always the new release is backward compatible with all previous versions. The following sections summarises the new features which have been added to TANGO V9.

Enum Attributes

A common need in control systems is to allow only a restricted set of values for certain attributes e.g. SINGLE, MULTI, HYBRID filling patterns. This is usually solved in programming languages using enumerated types. This needs special treatment when traversing the network. Version 9 has implemented enumerated attributes to make programmers and users lives easier.

Forwarded Attributes

Attributes are the well defined data exchanged between device servers and clients. In a hierarchical distributed control system system it is common to build hierarchies of device servers. When doing so it is very useful to be able to make some of the attributes of sub-devices available as attributes of other higher level devices. In the past this could only be done by manually re-programming the passing of the attribute of the sub-device in the higher level device. Version 9 has imported the notion of forwarded attributes which require no programming. They only need to be declared in the database. Less code means more reliability.

Pipes

TANGO supports a wide set of data types. For exchanging data these are mostly made up of scalars or arrays of scalars of the same type. This is sufficient in most cases. However there are some cases which have a need to exchange data of varying types coming from the same source. A typical example of this is when returning scanned data from an experiment scanning server. In this case it is useful to mix data types on the same communication channel. A new data communication channel, called pipes, has been introduced to address this problem. Version 9 introduces Tango Pipes for sending and receiving blobs of data of mixed types. Blobs of blobs are supported too.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

RECENT ADVANCEMENTS AND DEPLOYMENTS OF EPICS VERSION 4

G. White, M. Shankar, SLAC, Menlo Park, California, USA
A.N. Johnson, S. Veseli, ANL, Argonne, Ilinois, USA
A. Arkilic, L.B. Dalesio, M. Davidsaver, M.R. Kraimer, N. Malitsky, B.S. Martins, BNL, Upton, Long Island, New York, USA
M. Sekoranja, Cosylab, Ljubljana, USA
D.G. Hickin, DLS, Oxfordshire, England
T. Korhonen, ESS, Lund, Sweden
G. Shen, FRIB, East Lansing, Michigan, USA
R. Lange, ITER Organization, St. Paul lez Durance, France
S.M. Hartman, K. Kasemir, ORNL, Oak Ridge, Tennessee, USA

Abstract

EPICS version 4 is a set of software modules that add to the base of the EPICS toolkit for advanced control systems. Version 4 adds the possibility of process variable values of structured data, an introspection interface for dynamic typing plus some standard types, high-performance streaming, and a new front-end processing database for managing complex data I/O. A synchronous RPC-style facility has also been added so that the EPICS environment supports service-oriented architecture. We introduce EPICS and the new features of version 4. Then we describe selected deployments, particularly for high-throughput experiment data transport, experiment data management, beam dynamics and infrastructure data.

EPICS BASE AND EPICS VERSION 4

The Experimental Physics and Industrial Control System (EPICS), is a software framework for high-performance distributed control and data acquisition in large scientific instruments, such as accelerators and telescopes.

The "base" of EPICS is software for supervisory control, closed loop control, archiving, alarm management, timing and other aspects of front-end processors and device facing hardware. Typically hosted in an embedded system processor such as RTEMS or VME, this software and its host are collectively known as the IOC (Input/Output Controller) in an EPICS control system. IOCs are optimized for low-latency I/O. They control and/or monitor a collection of devices such as actuators and measurement diagnostics. Each IOC node contains a memory resident real-time database.

The IOC database is a set of "smart" records, which are interconnected in a *data flow* pattern. They're smart in that their field values may come directly from hardware, or a result of processing that was dependent on the type of record. The records may contain "device support" code, to interface the processing to physical devices through device drivers. Much more information can be found on the EPICS base at [1, 2].

EPICS base and the software extensions built on top of it have proven very successful for the control aspects of scientific instruments, providing excellent low level I/O, DAQ and optimal control.

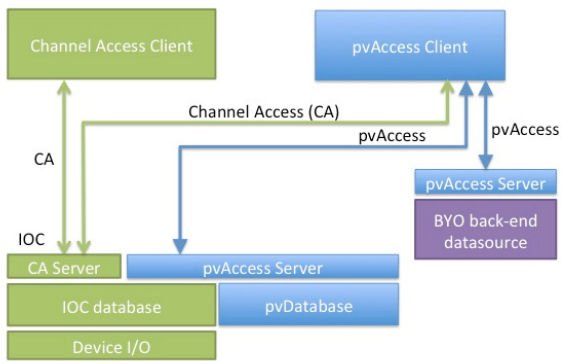


Figure 1: The basic architecture of EPICS version 4, showing the classic "base" components of EPICS in green, and components added by version 4 in blue. A new protocol, pvAccess, transports potentially complex data types encoded by pvData. pvAccess connects clients to IOCs, which may include a new database type, pvDatabase to support complex data processing services (and is the basis of the high-performance detector streaming applications described here), and to middleware services that can themselves connect to enterprise data stores, the web, etc. Note that hardware device I/O remains under the IOC database.

EPICS Version 4 for New Controls Problems

Modern control systems of large instruments call for more science to be done in the control system itself, than is possible with the EPICS base software alone. High output detectors call for pipelined data processing; physics applications deal with systems of process variables and their values rather than one process variable (PV) at a time; process variables may be structured, multi-modal, or require significant metadata associated with values in order to be understandable (such as images, or timing related metadata to give context when reviewing archives).

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

TOWARDS BUILDING REUSABILITY IN CONTROL SYSTEMS – A JOURNEY

Puneet Patwari, Amar Banerjee, G. Muralikrishna, N. Swaminathan, Subhrojyoti Roy Chaudhuri
Tata Research Development and Design Centre, Pune 411013, India

Abstract

Development of similar systems leads to a strong motivation for reuse. Our involvement with three large experimental physics facilities led us to appreciate this better in the context of development of their respective monitoring and control (M&C) software. We realized that the approach to allowing reuse follows the onion skin model that is, building reusability in each layer in the solution to the problem. The same motivation led us to create a generic M&C architecture through our first collaborative effort which resulted into a fairly formal M&C domain model. The second collaboration showed us the need to have a common vocabulary that could be used across multiple systems to specify respective domain specific M&C solutions at higher levels of abstraction implemented using the generic underlying M&C engine. This resulted in our definition and creation of a domain specific language for M&C. The third collaboration leads us to imagine capturing domain knowledge using the common vocabulary which will substantially further reuse, this thought is already demonstrated through a preliminary prototype. We discuss our learnings through this journey in this paper.

INTRODUCTION

Monitoring and control systems are central to the working of projects such as SKA[1], ITER [2] and so on. These projects incorporate wide variety of heterogeneous systems and subsystems which require supervisory controllers for coordination. Our involvement with these projects gave us opportunity to learn and understand the kind of challenges involved in building monitoring and control solutions for such systems. One of our key observations is that these projects do reuse a lot of artefacts for the purpose of final implementation of their control systems. However, they still incur a huge amount of cost due to the effort they spend in the initial phases of the development life cycle. We noticed that this effort could also be substantially reduced since it showed large commonality in the type of activities taking place in each phase of their development life cycle.

Motivated by this observation, we started to analyze the prospect of a generic M&C architecture. This led to the creation of a generic M&C design and a prototype to demonstrate it in the context of GMRT. The design was inspired by the data driven paradigm and resulted in identifying a set of engines that could configure themselves based on the supplied input data that described the problem context. This approach enabled capturing the abstract model behind this input data eventually serving as the generic domain or specification model to capture the details of any M&C problem. Our first implementation realized parts of this model based on the format of the underlying execution engine which resulted into fragmentation and duplication of the M&C problem spec. This showed us the need for an integrated environment which could ensure integrity and consistency in the M&C problem specification. We recognized the need for a domain specific language (DSL) [3] to enable specification of any M&C problem so that the solutions created using the DSL could be analyzed independently. Our DSL work showed us the need for an environment which could be made aware of the application domain through its support for extensibility, analyzability, re-targetability and so on. We realized that such an environment would enable reusing a lot of domain knowledge which would enhance consistency in the entire M&C development process.

In this paper we start with a discussion on the current practice and challenges that motivated our research. Next, we highlight the proposed solutions adapted throughout our journey. Subsequently, we provide a view of our current implementation followed by the section which summarizes and concludes the paper with a futuristic view.

STANDARD PRACTICE AND CHALLENGES

Most projects start working on the requirement and design of their M&C systems from the scratch. As a result each project or groups within a project end up creating their own version of the concepts around a general problem domain such as M&C. This leads to some re-invention of concepts that are already created in another project. This point towards the lack of reusable artefacts except implementation packages across a problem domain that could enhance reusability in the entire development life cycle.

System engineering language such as SysML [4][5]provides a convenient way of expressing the designs in most of the projects. However, since much of the M&C concepts are not built into the vocabulary of SysML, it is common for different groups within projects to define the M&C vocabulary independently using SysML. Unfortunately such definitions are mostly not shared across groups. This leads to non-uniformity in the definition and usage of the M&C concepts across groups within projects. Hence it requires manual effort to

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

OPEN SOURCE CONTRIBUTIONS AND USING OSGI BUNDLES AT DIAMOND LIGHT SOURCE

M. Gerring, A. Ashton, R. Walton, Diamond Light Source, Oxfordshire, UK

Abstract

This paper presents the involvement of Diamond Light Source (DLS) with the open source community, the Eclipse Science Working Group and how DLS is changing to share software development effort better between groups. The paper explains moving from product-based to bundle-based software development process which lowers reinvention, increases reuse and reduces software development and support costs. This paper details specific ways in which DLS are engaging with the open source community and changing the way that research institutions deliver open source code.

INTRODUCTION

Diamond Light Source [1] is a third-generation 3 GeV synchrotron light source based on a 24-cell double-bend achromatic lattice of 561m circumference. The photon output is optimised for high brightness from undulators and high flux from multi-pole wigglers. The accelerators and first phase of seven photon beamline were constructed from 2002 to 2007; a second phase of fifteen photon beamlines from 2006 to 2012; and a third phase of ten photon beamlines was approved in 2011 with construction due to finish in 2017-8.

As well as the construction of the synchrotron, the early phases of the project saw choices about the software which would be deployed on site. For hardware control such as motors, the EPICS framework was chosen which included a data driven user interface called EDM [2] for configuring devices. The acquisition and online data analysis system was developed from a product in operation at the SRS [3] called GDA [4] previously presented at ICALEPCS.

Diamond Light Source (DLS) have switched GDA (client) and a standalone analysis product called DAWN [5] to load using a system called OSGi (Open Service Gateway Initiative) [6]. In addition, the EPICS/EDM screens are planned to be phased out of active support [7] in line with the move to RHEL7. The next generation of software for controls and acquisition is based on Eclipse Rich Client Platform (RCP). This allows software products to be built from OSGi bundles and features to be developed which are interoperable between controls, acquisition and data analysis software. This paper details how DLS interacts with open source technology to deliver feature rich, interoperable and reusable software systems across groups and in the wider community.

OPEN SOURCE 'ECOSYSTEM'

Eclipse RCP [8] is a software technology which has been available for more than a decade. It is used to build user interface applications and OSGi servers. DLS are utilizing it as a platform to deliver native user interface clients. One of the first RCP clients in production for acquisition was on the B18 beamline, presented at ICALEPCS 2011 [9].

The Eclipse Foundation is also an open source publisher which verifies open source code for intellectual property (IP) and software license. Its rigorous IP process [10] renders code safe for commercial companies and institutions to reuse at reduced risk of litigation. It provides a rich open source feature set, similar to that on offer by the Apache Foundation. This 'ecosystem' of IP checked bundles has provided many useful features which DLS have been able to reuse within software products.

In 2014 DLS proposed an Eclipse project [11] to make aspects of its DAWN product open source and IP checked by the Eclipse Foundation. This project was granted and IP checking is active, nearing completion at the time of this conference.

How Open Source Works

The procedure in scientific institutions active in open source software release has often been to provide source using GNU Public License (GPL). Various mechanisms have been used to do this for example: zip file on an ftp site, by implementing a web site using a technology such as Redmine or by using an open repository site like github. These approaches are now considered unsafe however because the GPL can force institutions to release previously unready or non-public code. More importantly, the source code has often not been IP checked before it is released. Foundations such as Apache and Eclipse provide an IP checking service. The service gives confidence to the copyright holder that they have provided something for which they have a lower risk of litigation. The source code of the software is also more likely safe to be re-used. This promotes wider use and contributions from outside collaborators.

ECLIPSE SCIENCE WORKING GROUP

DLS, Oak Ridge National Laboratory [12] and IBM have formed an Eclipse Science Working Group (SWG) in conjunction with ten other members ranging from small contractors to large commercial companies. The SWG members have started projects such as Triquetrum, Chemclipse and Integrated Computational Environment (ICE), links to which can be found on the web site [13].

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

QUASAR - A GENERIC FRAMEWORK FOR RAPID DEVELOPMENT OF OPC UA SERVERS

P. P. Nikiel, B. Farnham, S. Schlenker, C.-V. Soare, CERN, Geneva, Switzerland
V. Filimonov, PNPI, Gatchina, Russia
D. Abalo Miron, University of Oviedo, Spain

Abstract

This paper describes a new approach for generic design and efficient development of OPC UA servers. Development starts with creation of a design file, in XML format, describing an object-oriented information model of the target system or device. Using this model, the framework generates an executable OPC UA server application, which exposes the per-design OPC UA address space, without the developer writing a single line of code. Furthermore, the framework generates skeleton code into which the developer adds the required target device/system integration logic. This approach allows both developers unfamiliar with the OPC UA standard, and advanced OPC UA developers, to create servers for the systems they are experts in while greatly reducing design and development effort as compared to developments based purely on COTS OPC UA toolkits. Higher level software may further benefit from the explicit device model by using the XML design description as the basis for generating client connectivity configuration and server data representation. Moreover, having the XML design description at hand facilitates automatic generation of validation tools. In this contribution, the concept and implementation of this framework named `quasar` (acronym for **q**uick OPC **UA** **s**erver gener**a**tion f**r**amework) is detailed along with examples of actual production-level usage in the detector control system of the ATLAS experiment at CERN and beyond.

INTRODUCTION AND MOTIVATION

Distributed control systems require middleware – software which transfers data between system components. The ATLAS Detector Control System (DCS) [1] is an example of such a distributed control system, organized as a hierarchical mesh of heterogeneous components. The middleware must be capable of handling numerous data models, while being portable and performant at the same time. For the ATLAS DCS, OPC Unified Architecture (further on: UA) [2] has been selected as its new standard middleware [3] for device integration mainly due to its object oriented design and platform independence. A common approach to create UA servers for the various device types allows to reduce development and maintenance costs.

Apart from obvious common functionality in which identical software parts were identified (such as server startup code, logging implementation etc.), it became evident that development efforts could be largely reduced if the data model was considered a parameter of a generalized UA server. Such a data model, augmented with additional information, is subsequently called *design*. If the format of the design is sufficiently rich to describe and model (potentially complex) subsystems, big parts of an UA server implementation may be automatically created (generated). Thereafter, hand-written custom code is only necessary for providing high level 'business logic' between the generated parts and the handling of the specific subsystem type (e.g. a hardware access library or protocol implementation). Such hand-written code may be very complex depending on its functionality requirements. We chose to call this code *device logic*.

In the following sections we explain the approach of generating UA servers from the preparation of a server design up to obtaining a functional application.

QUASAR ARCHITECTURE

Figure 1 gives an overview of the different layers of `quasar` put into context. Controllable devices or systems are accessed using their specific access layer – often provided together with the specific device. The device logic layer functions as interface with the high level layers of `quasar` which comes in several modules covering different functionality aspects. The address space module lies on the UA end of the server, exposing data towards UA clients, and is implemented using a commercial UA SDK [4]. A configuration module facilitates address space and device instantiation and the definition of their relations. XML is used as configuration format backed by XML schema definitions. A XML schema to C++ mapping generator (here: `xsd-cxx`) is used to build actual instances from configuration files. An additional subsystem called 'calculated items', operating entirely in the address space, enables creation of new variables which are derived from existing ones using mathematical functions. `quasar` comes further with optional modules such as component based logging, certificate handling, server metadata and embedded python processing.

MODELLING DEVICES OR SYSTEMS

In the generic approach of `quasar`, an object oriented model was chosen for two reasons: object orientation is well known and widely understood, and UA itself follows the object orientated paradigm. The purpose of modelling is to establish a comprehensive device or protocol characterization using classes, variables, methods and the relations between them. Classes are types of particular objects. Variables belong to classes and are factual vectors of data while class methods process the associated data. The purpose of relations is to model aggregations and type hierarchies.

Once the model is prepared, it has to be codified in a common format – we call this the *design file*. `quasar` uses the

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

DISRUPTOR - USING HIGH PERFORMANCE, LOW LATENCY TECHNOLOGY IN THE CERN CONTROL SYSTEM

M. Gabriel, R. Gorbonosov, CERN, Geneva, Switzerland

Abstract

Accelerator control systems process thousands of concurrent events per second, which adds complexity to their implementation. The Disruptor library provides an innovative single-threaded approach, which combines high performance event processing with a simplified software design, implementation and maintenance. This open-source library was originally developed by a financial company to build a low latency trading exchange. In 2014 the high-level control system for CERN experimental areas (CESAR) was renovated. CESAR calculates the states of thousands of devices by processing more than 2500 asynchronous event streams. The Disruptor was used as an event-processing engine. This allowed the code to be greatly simplified by removing the concurrency concerns. This paper discusses the benefits of the programming model encouraged by the Disruptor (simplification of the code base, performance, determinism), the design challenges faced while integrating the Disruptor into CESAR as well as the limitations it implies on the architecture.

INTRODUCTION

CESAR is the high level software used to control CERN experimental areas. The experimental areas are composed of eleven beam lines used by experimental physicists for fixed target research and detectors tests. Four beam lines are located in the East Area, using a beam extracted from the PS ring, and seven are in the North Area, using a beam extracted from the SPS ring. The core of CESAR is responsible for the data acquisition of all the devices controlling these beam lines. While refactoring this data acquisition part of Cesar, we decided to use the Disruptor library in order to simplify the design of the code handling the 2500 asynchronous event streams coming from these devices.

In the last decade, the actors of the world of finance and high frequency trading have been involved in an arms race to build exchanges and trading robots that can operate at the nanosecond scale. From time to time, some technologies created by the massive investments in this field are shared with the community [1].The Disruptor library was created by LMAX [2] -a London-based financial company- in order to develop a low-latency forex [3] trading venue [4]. In the early design phase, they tried different approaches: functional programming, Actors, SEDA [5], and noticed that they could not achieve the required latency because the cost of queuing was higher than the time spent executing the business logic. They finally settled on an innovative design and decided to open source it.

THEORETICAL BACKGROUND

The base idea around the Disruptor is to make the most of the available CPU resources, following a concept that its creators call ‘mechanical sympathy’. This term coming from the car racing world is used to describe software working in harmony with the hardware design, similar to a driver understanding how a car works in order to achieve the best performance. Since the appearance of multicore CPUs, we have heard expressions like “the free lunch is over” [6] and there is a general belief that CPUs are not getting any faster. Although their clock speed is not getting higher, modern CPUs have brought significant performance improvements. Regrettably, the progresses made in hardware are often lost by software designs that do not consider how modern processors work.

Feeding the Core

The most important aspect to consider in order to use a processor efficiently is to feed it correctly. Processors are very fast, but this speed is of little use if they spend most of their time waiting for the data they need to process. Looking at a simplified view of the memory architecture of a modern CPU such as Intel’s Sandy Bridge (Fig. 1), we see that the cores read data from several cache layers (L1, L2, and L3).

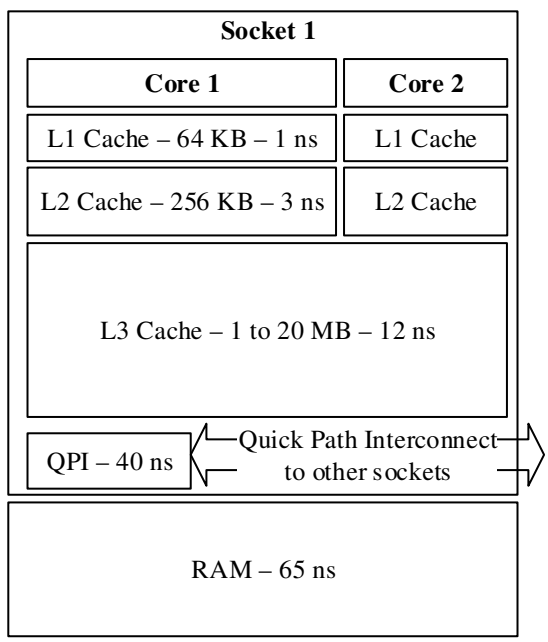


Figure 1: Memory Hierarchy

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

ACCELERATOR MODELLING AND MESSAGE LOGGING WITH ZeroMQ

J. Chrin, M. Aiba, A. Rawat, Z. Wang, Paul Scherrer Institut, 5232 Villigen PSI, Switzerland

Abstract

ZeroMQ is an emerging message-oriented architecture that is being increasingly adopted in the software engineering of distributed control and data acquisition systems within the accelerator community. The rich array of built-in core messaging patterns may, however, be equally applied to within the domain of high-level applications where a seamless integration of accelerator models and message logging capabilities, respectively, serve to extend the effectiveness of beam dynamics applications and allow for their monitoring. Various advanced patterns that include intermediaries and proxies further provide for reliable service-oriented brokers, as may be required in real-world operations. A report on an investigation into ZeroMQ's suitability for integrating key distributed components into high-level applications, and the experience gained, are presented.

MOTIVATION

ZeroMQ [1] is an emerging message-oriented architecture that has already made a profound impact within the wider accelerator and experimental physics community. A favourable evaluation from among several contemporaries [2] has since seen it target existing CORBA (Common Object Request Broker Architecture) [3] systems at several facilities [4–6]. It is also being adopted in message-based data acquisition systems [7–10] and has provided the means by which beam synchronous data is retrieved in high-frequency pulsed accelerators [9, 11]. Its apparent identification as today's middleware of choice has triggered an interest for its applicability within the domain of high-level applications at SwissFEL, Switzerland's X-ray Free-Electron Laser Facility [12, 13]; where CORBA may once have been used at a previous facility for incorporating distributed components into beam dynamics applications [14, 15], ZeroMQ presents itself as a viable, state-of-the-art, alternative that deserves consideration. Of particular interest is the integration of accelerator models and message logging capabilities which respectively serve to extend the effectiveness of beam dynamics applications and allow for their monitoring.

DISTRIBUTED COMPUTING WITH ZeroMQ

ZeroMQ is a lightweight, socket-like, asynchronous messaging library that provides for the transport of raw message buffers in a flexible and scalable distributed computing environment. The idiosyncratic name lends itself to the project's ambition to reach maximal performance by minimizing latency, copying, and the necessity for brokers (i.e. their numbers approach the limit of *Zero*). A first investigation into the ZeroMQ library already reveals a number of compelling features:

- A rich array of messaging patterns, including the familiar request-reply, publish-subscribe and push-pull (pipeline) patterns, each of which defines a distinct network topology.
- The availability of both unicast (inproc, ipc, tcp) and multicast (pgm, epgm) transport layers.
- The ability to use these patterns and transports as building blocks to establish connections between processes, with or without intermediate brokers/proxies.
- Support for multipart messages, which allow multiple frames to be concatenated into a single message to be sent over the network.

That these features are all available in a *single* library is positively favourable. (By comparison, CORBA requires separate libraries for their event driven and other services.) Furthermore, an active ZeroMQ community provides support for numerous platforms and an increasing multitude of programming languages.

Despite these benefits, some important components, that are outside ZeroMQ's stated interest, still need to be catered for to attain a fully fledged distributed infrastructure:

- A Name Service that translates logical addresses into bind/connect endpoints.
- An Implementation Repository for the activation and re-activation of servers.
- Support for object serialization.

With these shortcomings identified, what then are the remedies? Although a name resolution service may be developed from among ZeroMQ's architectural patterns, for the present time, the use of configuration files for publicising tcp/ip addresses is manageable. Most CORBA developers will have become accustomed to an Implementation Repository that interacted with the server's Object Adapter Mediator to handle the administrative aspects of server (re-)activations [14]. A similar setup for ZeroMQ would ensure that applications are never starved of the services that they require. The lack of an interface to serialize structured data may, at first, appear as a glaring omission given that most use cases would require it. Fortunately, this situation is redeemed through third-party solutions which vary in form, complexity and performance, endowing developers with the prerogative to choose that which best suits their needs.

Having now gained an insight into ZeroMQ's capabilities, its applicability to within the beam dynamics environment is readily recognized. The request-reply messaging pattern,

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

WHY SEMANTICS MATTER: A DEMONSTRATION ON KNOWLEDGE-BASED CONTROL SYSTEM DESIGN

W. Pessemier, G. Raskin, and H. Van Winckel, Institute of Astronomy, KU Leuven, Leuven, Belgium
P. Saey and G. Deconinck, ESAT, KU Leuven, Leuven, Belgium

Abstract

Knowledge representation and reasoning are hot topics in academics and industry today, as they are enabling technologies for building more complex and intelligent future systems. At the Mercator Telescope, we've built a software framework based on these technologies to support the design of our control systems. At the heart of the framework is a metamodel: a set of ontologies based on the formal semantics of the Web Ontology Language (OWL), to provide meaningful reusable building blocks. Those building blocks are instantiated in the models of our control systems, via a Domain Specific Language (DSL). The metamodels and models jointly form a knowledge base, i.e. an integrated model that can be viewed from different perspectives, or processed by an inference engine for model verification purposes. In this paper we present a tool called OntoManager, which demonstrates the added value of semantic modeling to the engineering process. By querying the integrated model, our web-based tool is able to generate systems engineering views, verification test reports, graphical software models, PLCopen compliant software code, Python client-side code, and much more, in a user-friendly way.

INTRODUCTION

Semantic models consist of pieces of information, and the relationships between those pieces. The ability to link any piece of information with another, thereby conveying the meaning (semantics) of the information, is what sets them apart from more "rigid" models such as those found in relational databases or object-oriented software. Semantic models are therefore well suited to represent all sorts of knowledge about the real world. At the Belgian Mercator Telescope (La Palma, Spain) we use the expressive power of semantic models to capture engineering knowledge of the telescope control system, which is being ported to a Programmable Logic Controller (PLC). As will be demonstrated in this paper, we have developed systems, electrical and software engineering models of several subsystems of the telescope, and successfully used these models for documentation, verification and implementation purposes.

FRAMEWORK ARCHITECTURE

As shown in Fig. 1, we have built a software framework centered around a Knowledge Base (KB) that combines (integrates) metamodels and models [1]. This KB can be queried by a tool which we developed (OntoManager), and the results of those queries can be fed into a template system to produce documents such as web pages and source code files. The metamodels provide the building blocks to construct models.

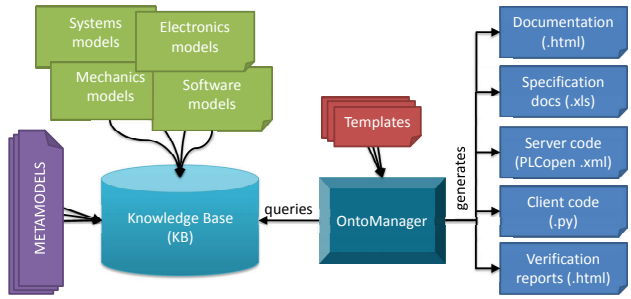


Figure 1: Framework architecture.

They are true *ontologies* as they define a vocabulary in such a way that the meaning of this vocabulary is well defined. Fig. 2 shows small excerpts of our systems metamodel (sys), mechanics metamodel (mech) and electronics metamodel (elec). They define concepts ("classes") such as `sys:Feature` and `mech:Assembly`, and relationships ("properties") such as `sys:hasFeature` and `elec:isConnectedTo`. Unlike a simple vocabulary, the meaning of these terms is further constrained whenever possible. For instance, the definition of a `mech:Assembly` says that "something" is a mechanical assembly if and only if it has at least two mechanical parts. Models based on this vocabulary will have to adhere to these constraints in order to be valid. For the underlying formal semantics (SubClassOf, EquivalentTo, Domain, ...) we depend on the Semantic Web standards RDFS (Resource Description Framework Schema) and OWL (Web Ontology Language) [2].

Systems metamodel (*sys*)

```
ObjectProperty: sys:hasFeature
  Domain: sys:System
  Range: sys:Feature
Class: sys:Feature
Class: sys:System
  EquivalentTo: sys:hasFeature some sys:Feature
```

imports

Mechanics metamodel (*mech*)

```
Class: mech:Part
  SubClassOf: sys:Feature
Class: mech:Assembly
  EquivalentTo: sys:hasFeature min 2 mech:Part
```

imports

Electronics metamodel (*elec*)

```
Class: elec:Conductor
  SubClassOf: mech:Part
  EquivalentTo: elec:carries exactly 1 quantity:ElectricCharge
  spin:constraint:
    [ a       sp:Construct .
      sp:text "CONSTRUCT { _:b0 a          spin:ConstraintViolation ;
                           _:b0 rdfs:label 'Signal mismatch!' .      }
               WHERE     { ?this  elec:carries       ?thisSignal .
                           ?this  elec:isConnectedTo ?other .
                           ?other elec:carries       ?otherSignal .
                           FILTER(?thisSignal != ?otherSignal)       }" ]
Class: elec:Contact
  EquivalentTo: sys:Interface and elec:Conductor
ObjectProperty: elec:isConnectedTo
  Characteristics: Symmetric
  Domain: elec:Contact
  Range: elec:Contact
```

Figure 2: Small excerpts of some metamodels.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

TRIGGER AND RF DISTRIBUTION USING WHITE RABBIT

T. Włostowski, J. Serrano, G. Daniluk, M. Lipiński, CERN, Geneva, Switzerland
F. Vaga, University of Pavia

Abstract

White Rabbit is an extension of Ethernet which allows remote synchronization of nodes with jitters of around 10ps. The technology can be used for a variety of purposes. This paper presents a fixed-latency trigger distribution system for the study of instabilities in the LHC. Fixed latency is achieved by precisely time-stamping incoming triggers, notifying other nodes via an Ethernet broadcast containing these time stamps and having these nodes produce pulses at well-defined time offsets. The same system is used to distribute the 89us LHC revolution tick. This paper also describes current efforts for distributing multiple RF signals over a WR network, using a Distributed DDS (Direct Digital Synthesis) paradigm.

SYNCHRONIZATION IN WHITE RABBIT

Both described systems require a precise time and frequency reference, provided by a WR network and locked to a GPS receiver or an atomic clock. WR achieves its subnanosecond accuracy and low jitter by employing multiple techniques:

- Distribution of frequency reference encoded in the physical data stream (Synchronous Ethernet).
- Coarse clock offset measurement by timestamping Ethernet frames using the Precision Time Protocol (PTP, IEEE1588).
- Fine offset compensation by tracking the phase shift between the outgoing and incoming clock/data stream.

Further details on the synchronization algorithms used in WR can be found in [1].

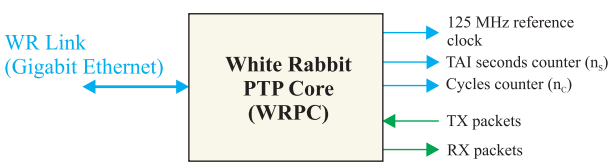


Figure 1: The White Rabbit PTP Core (WRPC).

The RF and Trigger distribution systems are interfaced with the WR network by the WR PTP Core (WRPC) [2], a standardized HDL block implementing the full WR synchronization stack. The WRPC is integrated in the FPGA firmware and delivers a 125 MHz reference frequency and time-of-day (counters of seconds n_S and 125 MHz clock cycles n_C). Furthermore, the WRPC can be used as general-purpose Ethernet Media Access Controller (MAC), sending and receiving Ethernet frames coming from the user's FPGA design (see Figure 1).

RF DISTRIBUTION

Introduction

The RF distribution system is part of a larger project which aims to merge the two historically separate timing systems:

- General purpose timing, usually referenced to the Universal Coordinated Time (UTC) timescale, which usually drives beam injection/extraction, synchronizes currents in the magnets and provides UTC-traceable timestamps for any sort of events that may occur in the machine.
- Beam-synchronous timing, using the bunch or RF frequency as the reference. Its natural applications are driving the RF cavities, synchronizing data acquisition from beam instrumentation or providing precise collision timestamps for event reconstruction (e.g. in the LHC Experiments).

The classic way of distributing a frequency reference in a timing system is to encode it in the data stream (using Manchester or 8B10B encoding) at the master node and recover it in the slave nodes using a Clock-Data Recovery (CDR) PLL. Examples of such timing systems can be found in [3] and [4]. This approach, however, can provide only one reference frequency per each physical link, necessitating the use of two separate timing networks. In many cases the machine's equipment needs simultaneous access to both timing systems, resulting in duplication of cabling and electronics.

Furthermore, many traditional systems experience troubles tracking large frequency changes (e.g. ramping of the RF in low energy or ion machines), due to bandwidth limitations of the CDR PLLs in the deserializer chips.

The method we developed does not rely on physical frequency encoding, but uses the common notion of time and frequency provided by WR to drive RF synthesis in each node of the system, as depicted in Figure 2. The master node keeps is local DDS phase-locked to the RF reference input and broadcasts the DDS tuning values calculated by its PLL over the WR network. The slaves simply feed the received data to their local DDS synthesizers. Since the reference clocks are identical, the DDS in the slave node produces an exact copy of the RF clock coming to the master node.

RF Encoding and Broadcasting

The RF encoding done by the master node is illustrated in Figure 3. The central element of the system is the DDS synthesizer, which produces a sinusoidal signal of arbitrarily controlled frequency and phase. The synthesizer we used is a custom-designed FPGA core attached to a high speed DAC (Digital to Analog Converter). It employs the classical frequency synthesis approach, using a phase accumulator

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

THE PHASE-LOCKED LOOP ALGORITHM OF THE FUNCTION GENERATION/CONTROLLER

M. Magrans de Abril, Q. King, R. Murillo Garcia, CERN, Geneva, Switzerland

Abstract

This paper describes the phase-locked loop algorithms that are used by the real-time power converter controllers at CERN. The algorithms allow the recovery of the machine time and events received by an embedded controller through WorldFIP or Ethernet-based fieldbuses. During normal operation, the algorithm provides less than 10 µs of time precision and 0.5 µs of clock jitter for the WorldFIP case, and less than 2.5 µs of time precision and 40 ns of clock jitter for the Ethernet case.

INTRODUCTION

It is widely known that the existence of measurement and computation delays and jitter affect the stability and performance of the control algorithms [1–3]. It is also well known that as the control system becomes larger (i.e. spatial extension or number of jointly controlled elements), the complexity of the control algorithm increases and its performance and stability degrades [4]. This is precisely the situation of the control of power converters at CERN. For example, the LHC project required the joint control of over 1700 converters across the 27 km of underground tunnels with a tracking error of less than five part per million (ppm). The large spatial extension requires a correspondingly large communication network to transport the information, which increases the delay, jitter, and packet loss. If there are different subsystems to be jointly controlled (e.g. main dipole and quadrupole magnets, simultaneous setting of the real-time orbit corrections, synchronisation of injection and ejection converters, etc.), then the control algorithm should also consider the relative delays between them.

In order to minimise the above effects, the devices that require real-time synchronisation at CERN used to employ a special purpose communication network that transports timing information with small jitter and known delay [5, 6]. The information necessary to control a device is therefore sent using two different networks. An Ethernet network (known as the technical network) transports the soft real-time data. A second network (known as timing network) transports the hard real-time data containing such things as the machine time, accelerator state and events.

Since the year 2000, the power converter controllers known as Function Generation Controllers (FGC) have deviated from this general architecture [7]. As shown in Fig. 1 the FGC uses the fieldbus as the mechanism to receive both soft and hard real-time data from the general purpose Linux computer (known as the front-end). That is, the control and the timing networks converge on a front-end, instead of being connected directly to the embedded controller. This change reduces the installation cost and the number of connection related incidents. In the FGC case, further reduction of development, installation, and maintenance costs have been achieved by using commercial off-the-shelf components and protocols for the fieldbuses. That is, WorldFIP on the FGC version 2 (FGC2), and FGC_Ether on the FGC version 3 (FGC3) [8].

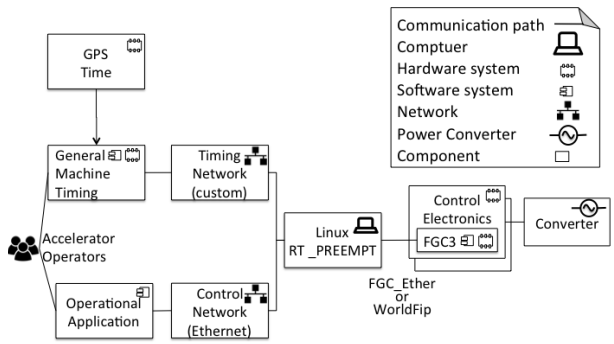


Figure 1: Transport of timing and control information from the GPS signal and the particle accelerator operator to the power converter using one fieldbus to transport timing, control, and monitoring information.

Removing the direct connection between the timing network and the FGC implies that the UTC time and its phase have to be recovered. The time with 20 ms precision is recovered by a periodic broadcast signal from the general purpose computer to the FGC, while the 20 millisecond phase is recovered using a Phase-Locked Loop (PLL) running on the FGC. This paper describes the requirements, design, and performance of the PLL algorithms used for both the FGC2 and FGC3.

SYNCHRONISATION OF THE LHC CONVERTERS

Requirements and Design

Between the year 2000 and 2003 the FGC2 was developed for the LHC. On one hand, the correct operation of the LHC required a radiation tolerant fieldbus and the synchronised application of the converter current across the accelerator with an accuracy of less than 5 ppm. This requires a time precision of 1 ms to achieve the synchronised operation of the converters [9]. It also required a jitter of less than 10 µs to have less than 1 ppm jitter induced noise from the sigma-delta ADC. On the other hand, it was not of primary concern the control and monitoring bandwidth between the operator and the controller because the circuits and the ramps in current are very slow. These requirements were fulfilled by a 2.5 Mbit/s WorldFIP fieldbus cycling at 50 Hz. Given the low jitter of this real-time fieldbus (i.e. $\sigma_j \approx 0.5\,\mu s$),

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

NEW EVENT TIMING SYSTEM FOR DAMPING RING AT SuperKEKB

H. Kaji, K. Furukawa, M. Iwasaki, T. Kobayashi, F. Miyahara, T.T. Nakamura,
M. Satoh, M. Suetake, M. Tobiyama, KEK, Japan
Y. Iitsuka, East Japan Institute of Technology, Japan
T. Kudou, S. Kusano, Mitsubishi Electric System & Service Co., Ltd, Japan
M. Liu, C. Yin, SINAP, China

Abstract

We utilize the Event Timing System to perform the operation of injector for the SuperKEKB accelerators. We construct the Sub-timing Station to manage two kinds of triggers for the Damping Ring devices. We developed new Event modules, VME-EVO and VME-EVE, to deliver triggers to the beam monitors at Damping Ring. The triggers for 84 beam position monitors can be provided with only 5 new Event modules. The timing accuracy of outputs from the new Event modules is ~15 ps. It well satisfies the requirements from the beam monitors at Damping Ring.

INTRODUCTION

SuperKEKB [1, 2] is the electron-positron collider which targets the world's largest record of luminosity. This is the upgrade machine of KEKB [3, 4] and one of the most important projects at KEK. SuperKEKB is designed to achieve 40 times larger luminosity than that of KEKB – it is realized from two times larger storage current and 20 times smaller vertical emittance.

The control system of injector linac (LINAC) [5] is extremely important for both KEKB and coming SuperKEKB. One of key technologies for the LINAC control is Pulse-to-Pulse Modulation (PPM). This injector "simultaneously" performs the top-up filling operations into four independent rings – two rings of SuperKEKB and two light sources [6, 7]. Therefore more than 150 of parameters to operate the LINAC beamline should be changed, pulse-by-pulse, typically in 50 Hz.

The Event Timing System of LINAC [8–10] works for not just the delivery of timing-triggers to the beamline devices but the above-mentioned PPM operation. It forms star topology optical network. The Event Generator (EVG) [11] is placed at Main Timing Station[1] – it is the focal point of optical network. The EVG controls the LINAC devices by delivering various kinds of Events. The Event Receivers (EVRs) [11] are placed at the individual edges of the optical network. The EVRs work for the beamline devices in the two kinds of ways when they receive the Events. One is the launching of triggers, like NIM or TTL, for indicating the timing of injection process. The other is the CPU interruption to switch the parameters of devices for the PPM.

We develop Sub-timing Station on one edge of Event network for the newly constructed Damping Ring (DR). The several Event modules including the newly developed modules are installed at DR Sub-timing Station since the delivery of timing-triggers towards the DR devices during the storage of beam-pulses is complicated.

In this report, we introduce the timing system at DR with the newly developed Event modules.

[1] We call it as Main Trigger Station in the KEKB era.

TRIGGER REQUIREMENTS AT DR

In this section, we explain the requirements to the timing system at DR. We need the special treatments to the triggers for the beam monitors since they are not synchronized with the LINAC operation. We explain the triggers for the beam position monitors (BPMs) as an example.

Operation Scheme of DR

DR is constructed for SuperKEKB. It is utilized in the injection of positron-pulses at LINAC. The frequency of RF cavity at DR is 508.89 MHz and it is same as that at MR. The harmonic number is decided to be 230 for making Bucket Selection of Main Ring (MR) [12, 13] efficient. The positron-pulses are stored at least 40 ms for damping their emittance.

In each positron injection, two of RF-buckets are occupied by the beam-pulses. The RF-buckets to be occupied are changed depending on the Bucket Selection in pulse-by-pulse.

Triggers for BPM

There are 84 BPMs at DR. They work during the storage of positron-pulses and measure the orbit. The timing of trigger for BPM is set by the delay value from the DR revolution as a reference and is independent from the LINAC operation.

The delay value is determined from two components. The individual BPMs have their own delay which are related with the installed positions along the DR beamline. In addition to these individual delays, there is the common delay which is related with the RF-buckets occupied with beam-pulses. The common delay must be set in the 508.89 MHz step since it is the frequency of RF cavities at DR. It should be adjusted, in pulse-by-pulse, because of Bucket Selection.

To validate the certain performance of BPMs, the accuracy (jitter) of the total delay values should be < 2 ns.

Summary of Requirements to the Timing System

Here we summarized the requirements to the timing system for the DR operation.

There are the triggers only for the DR devices like the beam monitors. They are delivered during the storage of

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

TIMING SYSTEM FOR THE HALPS/L3 ELI PROJECT

P. Camino, D. Monnier-Bourdin, Greenfield Technology, Bordeaux, France
M-A. Drouin, J.A. Naylon, B. Rus, FZU, Czech Republic
S. Telford, G. Johnson, C. Haefner, LLNL, Livermore, CA, USA

Abstract

The High Repetition-Rate Advanced Petawatt Laser System (HAPLS) forms part of the European Union's Extreme Light Infrastructure Beamlines project (ELI-Beamlines) which will be the first international laser research infrastructure of its kind. HAPLS will generate peak powers greater than one petawatt at a repetition rate of 10 Hz with 30fs wide pulses. ELI will enable unprecedented techniques for many diverse areas of research. HAPLS requires a high-precision timing system that operates either independently or synchronized with ELI's system. Greenfield Technology, a producer of mature picosecond timing systems for several years, has been hired by LLNL to provide just such a timing system. It consists of a central Master Timing Generator (MTG) that generates and transmits serial data streams over an optical network that synchronizes local multi-channel delay generators which generate trigger pulses to a resolution of 1ps. The MTG is phase-locked to an external 80 MHz reference that ensures a jitter of less than 10ps rms. The various qualities and functions of this timing system are presented including the LabVIEW interface and precision phase locking to the 80MHz reference.

INTRODUCTION

HALPS is planned to become PW workhouse of the ELI-beamlines facility composed of 3 other laser systems. ELI-beamlines located in Czech Republic is designed to explore fundamental physics under extreme conditions. It will contain the world's most powerful lasers for use by the international scientific community. ELI-beamlines is one of three laser facilities currently under construction as part of the ELI project (European Union's Extreme Light Infrastructure): ELI Attosecond located in Hungary to investigate natural phenomena on ultra-short timescales and ELI Nuclear Physics located in Romania dedicated to the new field of photonuclear physics are the two remaining facilities.

HALPS's short-pulse laser will use titanium-doped sapphire as its amplification medium. This laser is designed to convert the energy from the pump laser to 30-joule, 30-femtosecond pulses for a peak power exceeding 1 petawatt. HALPS is key to firing at 10Hz for hours at a time. The aim of such a laser is to generate secondary sources of electromagnetic radiation and accelerate charged particles. It will contribute to the development of laser-driven fusion power plants as well [1].

The HALPS block diagram is depicted on Figure 1. The main parts are: a pump laser + frequency converter and a short pulse laser + compressor.

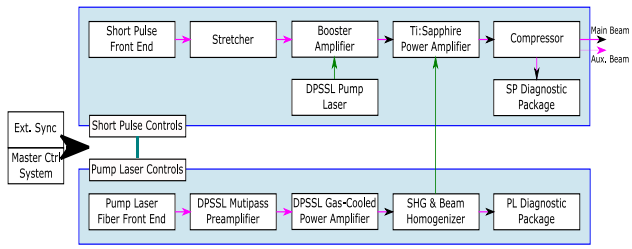


Figure 1 : HALPS block diagram.

We discuss the timing system required for such a laser.

PURPOSE

Synchronization of the pump laser beam and the short pulse laser beam requires generation of numerous precisely timed triggers. In addition the output of the laser needs to be synchronized in time with the balance of the ELI-Beamlines facility.

The purpose of the facility timing system is to generate and distribute over optical fibres a serially encoded, time critical data pattern to multiple zones. These zones contain multichannel delay generators that comprise the Local Timing System. The delay generators decode the data pattern contents and generate the electrical and optical triggers required by HALPS.

The low-jitter high precision timing system from Greenfield Technology has been selected by LLNL.

LOW JITTER – 1PS ELECTRICAL TIMING SYSTEM

Greenfield Technology has developed high precision timing systems for 10 years [2]. Its timing systems are already in use in several big physic instrument systems such as SOLEIL [3], LMJ, APOLLON [4], LULI (France) or Jupiter, NIF (LLNL, USA).

The one presented here has been customized to fulfill LLNL/HALPS requirements (output standard, clock frequency and content of the serial data stream). One of the major system specifications is the MTG phase locked specification to the 80MHz external clock with an rms jitter from clock to optical data stream less than 10ps. The different parts of the designed timing system are described in the following sections.

Master Timing Generator

The primary function of the MTG is the generation and distribution of timing data. It is composed of the Master Timing Transmitter and a Control Interface.

The Control Interface function serves the dual purpose of providing internal computation and control within the MTG and as interface to an external computer: Ethernet protocol communication using TCP/IP.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

ERL TIME MANAGEMENT SYSTEM*

P. Kankiya†, T. Miller, B. Sheehy, Brookhaven National Laboratory, Upton, NY 11973, USA

Abstract

The Energy Recovery LINAC (ERL) at BNL is an R&D project. A timing system was developed in conjunction with other available timing systems in order to operate and synchronize instruments at the ERL. This paper describes the time management software which is responsible for automating the delay configuration based on beam power and instrument limitations, for maintaining beam operational parameters, and respond to machine protection system.

INTRODUCTION

The Energy Recovery LINAC at BNL is an R&D project and has been constructed in a blockhouse at BNL's large complex [1]. It is advantageous to setup a prototype machine on the same site which has long running accelerators such as RHIC because of access to resources and expertise. The challenging part is to meet project specific needs that are unable to utilise the existing infrastructure. In case of ERL's timing system it was partly possible to use RHIC style delay generator boards but certain features required programming new or buying commercial off-the-shelf (COTS) delay generators. The timing clock and events are distributed to all parts of ERL such that Laser and RF are always phase locked. This is the basis for synchronizing all aspects of ERL behaviour. The timing system is clocked by a 4.69 MHz clock optically derived from the photocathode laser; the laser is itself phase-locked to the 704 MHz SRF gun that serves as the master clock of the system. The laser clock is distributed to the timing system, installed about 200 ft away from Laser room in a VME rack. The laser-clocked VME system triggers some hardware systems directly, and others are triggered by low jitter (10s of psec) COTS delay generators (Stanford Research Systems DG645), which are triggered by the VME system. The DG645's offer more flexibility in certain functions (eg, burst, inhibit, gate shape, pre-scaling) than the native VME system, and this hybrid approach expands the capability and adaptability of the timing system. This assembly of various delay units and monitoring software together constitute the timing system of ERL and is credited for tasks such as coherent operation of measuring instruments, automatic delay calculation, supporting system testing, integrating several key hardware components, beam parameter limitation, inhibiting diagnostic systems and Laser beam production.

SYSTEM LAYOUT

The timing system currently comprises of two V202 delay modules installed in two different VME chassis and four DG645 boxes. A top level flow-diagram of timing signals is illustrated in Fig. 1. The inception of timing signals is at a V202 delay module. These modules are capable of decoding RHIC event link events [2]. A C++ class called Accelerator Device Object (ADO) [3] is created for controlling V202 boards. This ADO lets a user have control of each delay channel by providing fields for delay values, output pulse width, source of trigger etc. Using the ADO interface the board is configured to operate in external clock mode and is connected to 4.69 MHz clock generated by the photocathode laser. All output channels of 202 modules can be triggered externally, manually or at an event. The software representation of a DG645 is also realised by ADO. These units are all configured to be externally triggered and concurrently triggered by V202 channels with a maximum lag of 213ns lag (one clock cycle).

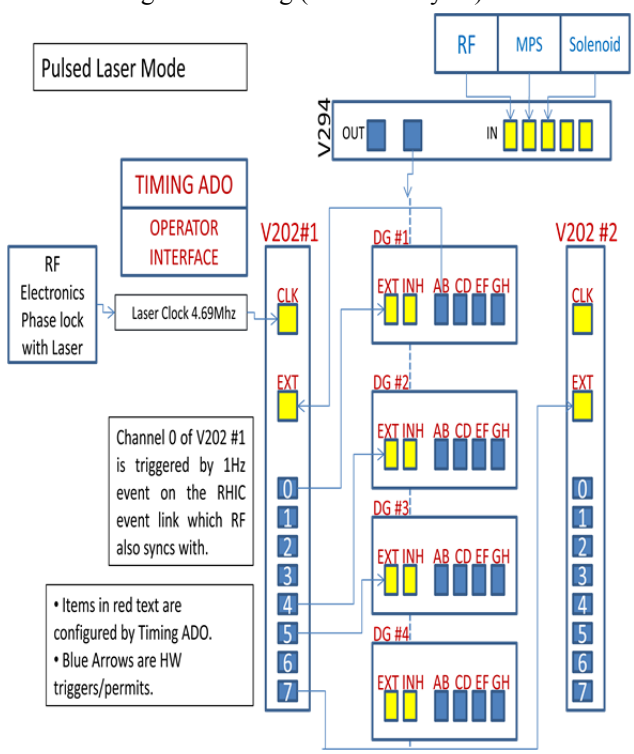


Figure 1: Block Diagram of the timing system with software and hardware components.

All configurable controls and statuses on these devices are displayed on user screens as ADO parameters. An ADO class referred to as master timing ADO is constructed to orchestrate timing signals across all ERL equipment via both styles of delay boxes such that the operator interface stays streamlined and unaffected by variety of hardware at lower level.

*Work supported by Brookhaven Science Associates, LLC under Contract No. DE-AC02-98CH10886 with the U.S. DOE.

† pkankiya@bnl.gov

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

MASSIVE: AN HPC COLLABORATION TO UNDERPIN SYNCHROTRON SCIENCE

Wojtek James Goscinski, Chris Hines, Paul McIntosh,
Monash University, Clayton, Australia
Keith Bambery, Ulrich Felzmann, Chris Hall, Anton Maksimenko,
Santosh Panjikar, David Paterson, Mark Tobin,
Australian Synchrotron, Clayton, Australia
Chris Ryan, Darren Thompson,
CSIRO, Clayton, Australia

Abstract

MASSIVE is the Australian specialised High Performance Computing facility for imaging and visualisation. The project is a collaboration between Monash University (lead), Australian Synchrotron (AS) and CSIRO, and it underpins a range of advanced instruments, including AS beamlines. This paper will report on the outcomes of the MASSIVE project since 2012, in particular focusing on instrument integration, and interactive access for analysis of synchrotron data. MASSIVE has developed a unique capability that supports an increasing number of researchers, including an instrument integration program to help facilities move data to an HPC environment and provide in-experiment data processing. This capability is best demonstrated at the AS Imaging and Medical Beamline (IMBL) where fast CT reconstruction and visualisation is essential to performing effective experiments. A workflow has been developed to integrate beamline allocations with HPC allocation providing visitors with access to a dedicated project space and a CT reconstruction service using a remote desktop environment. The work herein describes the processing and analysis workflows developed for processing AS data at the IMBL, Macromolecular crystallography beamlines (MX), X-ray Fluorescence Microscopy (XFM), and the Infrared Microspectroscopy Beamline (IR).

THE MASSIVE FACILITY

MASSIVE was established to provide data analysis and visualisation services to the imaging and instrumentation community. The facility provides computer hardware, software and expertise to drive research in the biomedical science, materials research, engineering, and neuroscience communities, and it underpins advanced imaging research in synchrotron X-ray and infrared imaging, functional and structural magnetic resonance imaging, X-ray computed tomography (CT), electron and optical microscopy. MASSIVE offers Australian scientists access to two specialized computing facilities at Monash University and AS, plus specialised services offered through the cloud.

MASSIVE underpins the processing and analysis requirements of beamlines at AS. The IMBL is able to produce data at over 500 Mbytes/s and data sets over a terabyte in size. This introduces obvious challenges for researchers to capture, process, analyze and visualise data in a timely and effective manner. Researchers are also increasingly eager to perform data analysis "in-experiment" to make quick decisions.

Hardware

MASSIVE provides two computers, M1 and M2, that operate at over 5 and 30 teraflops respectively, using CPU processing, and over 50 and 120 teraflops, respectively, using Nvidia and Intel co-processors.. The computers are connected using a dedicated network for fast file transfer and management*.

GPUs have proved an important part of the MASSIVE environment. A number of key applications, including the X-TRACT [1] CT reconstruction software, have been parallelized to take advantage of GPUs and optimized on MASSIVE which is critical to fast processing of data in near real-time at the IMBL. Furthermore, GPUs underpin interactive visualisation through the MASSIVE Desktop and through parallel rendering tools such as Paraview [2].

Both systems have a GPFS [3] file system that is capable of a combined 5GB+ per second write speed. Fast file system performance has proved essential to support both the fast capture of data from instruments, and file system intensive image processing.

Instrument Integration

MASSIVE runs a program for the integration of imaging instruments with high performance computing capability. This program helps develop workflows for scientists to process and visualise data as an experiment progresses or immediately after it completes. A number of instruments have been integrated with MASSIVE [4]. The IMBL CT reconstruction workflow (described within) is a major output of this program.

The MASSIVE Desktop

MASSIVE provides users with an easily accessible scientific desktop environment configured for analysis and visualisation of research data. It provides researchers with access to a range of existing tools and software, including commercial and open-source applications used for synchrotron data processing such as X-TRACT [1],

* Detailed specifications available at: https://www.massive.org.au/high-performance-computing/resources

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

DATABROKER: AN INTERFACE FOR NSLS-II DATA MANAGEMENT SYSTEM

A. Arkilic, L. Dalesio, D. Allan, D. Chabot Brookhaven National Laboratory, NSLSII
Upton, NY, USA

Abstract

A typical experiment involves not only the raw data from a detector, but also requires additional data from the beamline. This information is largely kept separated and manipulated individually, to date. A much more effective approach is to integrate these different data sources, and make these easily accessible to data analysis clients. NSLS-II data flow system contains multiple back-ends with varying data types. Leveraging the features of these (metadatastore, filestore, channel archiver, and Olog), this library provides users with the ability to access experimental data. The service acts as a single interface for time series, data attribute, frame data access and other experiment related information.

INTRODUCTION

The current and projected NSLS-II beamlines will support a wide range of disciplines from biology to physics with a wide range of techniques. Higher beam quality and new generations of area detectors are increasing data rates quickly. Increasingly all disciplines are using multi-element detectors to measure and record the results of the experiment. Current state-of-the-art detectors are capable of producing megapixel images with frame-rates in the kilohertz range which results in data rates of up to 800 Mb.s-1. Such data-rates are fully compatible with current off-the-shelf RAID storage technology, effectively allowing the data to be written to a storage medium as it is produced. The challenge for NSLS-II, however is to be able to process and analyse the data at this rate, with these volumes to allow the experimenter to not only make the "first-cut" for decision making during the experiment, but to retrieve data and analyse it for publication. NSLS-II middle layer beamline applications (Fig. 1), powered by MongoDB, the General Parallel File System (GPFS), and EPICS provide such capabilities.

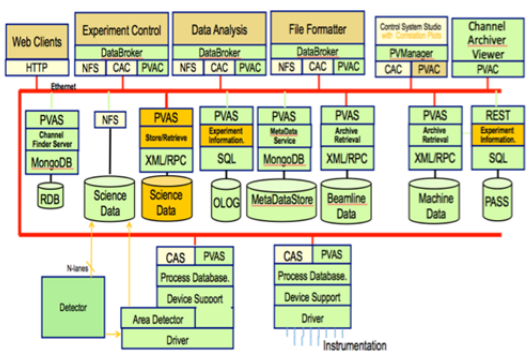


Figure 1: An overview of NSLS-II middle-layer applications.

Given the modular architecture and complexity of these various applications, a need for an easy point of access for all experimental data has occurred. In order to simplify the access to critical experimental data, we have developed databroker.

ARCHITECTURE

databroker is the pathway for data analysis tools to access data without any knowledge of experimental data. The library in production is implemented in Python and provides users with broker objects that contain Pandas data frames, allowing semi and un-structured experimental data to be accessed within scientific Python framework with ease via its narrow interface.

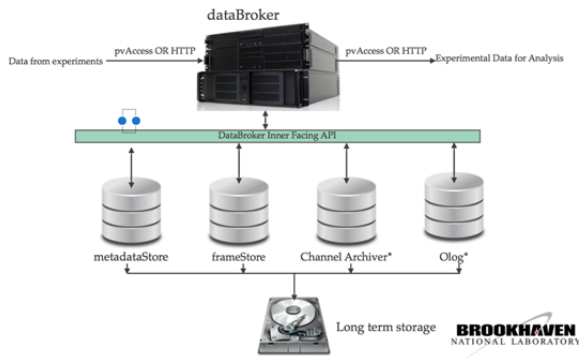


Figure 2: databroker Architecture Overview.

databroker uses the client APIs of metadatastore, filestore, channel archiver, and Olog (Fig. 2). Its primary goal is to provide users with run-based information. A run is defined as an end-to-end scan or sweep that is performed and captured either via NSLS-II run engine, BlueSky, or custom data acquisition framework that utilizes NSLS-II middle layer services. Run based data management is common to all beamlines and it provides users with the desired abstraction.

Deployment Strategy and Access Environment

All python clients that are part of NSLS-II middle layer get deployed to each beamline using the Anaconda python package management toolkit. The sandbox approach allows us to deploy packages that are required but not natively supported by the Debian 8. IPython notebook is also part of this deployment and it is the primary platform to access databroker. In addition to local access to databroker, NSLS-II beamlines provide access to an identical databroker session remotely using jupyter hub.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

MADOCA II DATA LOGGING SYSTEM USING NoSQL DATABASE FOR SPring-8

A. Yamashita*, M. Kago, SPring-8, Hyogo, Japan

Abstract

The data logging system for SPring-8 was upgraded to the new system using NoSQL database, as part of the MADOCA-II framework. The data logging system has been collecting the log data that is required for accelerator control without any issues since the upgrade. In the past, the MADOCA system powered by a relational database management system (RDBMS) had been operating since 1997. The MADOCA system had grown with the development of accelerators. However, the system with RDBMScould not handle new requirements like variable length data storage, data mining from large volume data and fast data acquisition. Therefore, new software technologies were developed to address these problems. In our proposed system, we adopted two NoSQL databases, Apache Cassandra and Redis, for data storage. Apache Cassandra is utilized for providing a perpetual archive. It is a scalable and highly available column oriented database suitable for time series data. Redis is used for the real time data cache because of a very fast in-memory key-value store. Data acquisition part of the new system was also built based on ZeroMQ message packed by MessagePack. The operation of the new system started in January 2015 after an evaluation term over one year.

INTRODUCTION

The new MADOCA-II data acquisition and storage system took over from the old system in January 2015. Since then, it has been under operation at the SPring-8 accelerator and beam lines control system without major troubles.

In the old MADOCA system [1], one large RDBMS [2] managed not only the equipment and operation parameters but also the log data [3]. It began the production run during the commissioning of SPring-8 in 1997 and has since grown as the accelerator has evolved. The number of signal data increased from 871 at the commissioning in 1997 to 27,626 at the end of 2014 and the size of the database expanded from 17.2GB (1998 one year) to 4 TB. With the progress of the accelerator and for the next generation SPring-8-II [4], we required more scalability, flexibility and maintainability for the database system especially for the log database. In The old MADOCA database system, almost every control application depended on a single database server and the server was required to be reliable and without down time. On the other hand, increase in the number of signals required more performance from one server. Therefore, we employed a fault tolerant database server and SAN (Storage Area Network) storages for reliability. Although, they are very reliable and have worked without trouble in years, they are expensive and pose a difficulty during scale up.

In the previous ICALEPCS, we reported the design and implementation of MADOCA-II database, which used NoSQL databases [5]. The system consisted of two databases, Apache Cassandra [6] for perpetual data store and Redis [7] for the real time data cache. The Apache Cassandra runs on a redundant cluster of commodity servers. It makes the cluster scale-out by simply adding additional server nodes. Data are replicated and distributed to nodes. In our case, up to two simultaneous node downs are tolerable. On the other hand, Cassandra scarifies the consistency of data. A Cassandra cluster distributes data to its nodes. It takes time (about 1 sec in our case) until every node has consistent data. We overcome the inconsistency of Cassandra by writing data into a very fast in-memory data cache; Redis. The data acquisition system writes data into Cassandra and Redis simultaneously. Two Redis servers are run in parallel for redundancy.

We installed the test system in SPring-8 accelerator and beam line control environment. The system acquires and stores the same set of data as MADOCA database system in parallel. During the one year test, we operated the database and data acquisition system with no major trouble. We obtained many know-hows during the test run. From January of 2015, we replaced old data acquisition and database system with new MADOCA-II system with some modifications based on know-hows that were obtained during the test run.

IMPLEMENTATION

Figure 1 shows the entire data acquisition and database system. Our previous paper [8] explained the data acquisition system. We did not make major modifications to it. The data acquisition system operates with the old system in parallel.

We developed a data sender running embedded computer as well as libraries for data reading, web system and alarm system running on client computers.

Data Senders

We developed two types data sender one is called po2m2db which runs on an embedded computer system. It extracts data from the shared memory inside, produces messages, and transmits them to a relay server asynchronously. The other data sender is called cc2m2db for an embedded system that has not enough resources to run po2m2db. The old data acquisition system [9], which runs on a workstations, outputs data to stdout. cc2m2db captures the stdout data, parse them, generates messages and send to the relay server.

* aki@spring8.or.jp

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

HDB++: A NEW ARCHIVING SYSTEM FOR TANGO

L. Pivetta, C. Scafuri, G. Scalamera, G. Strangolino, L. Zambon,
Elettra Sincrotrone Trieste, Trieste, Italy
R. Bourtembourg, J.L. Pons, P. Verdier, ESRF, Grenoble, France

Abstract

The TANGO release 8 led to several enhancements, including the adoption of the ZeroMQ library for faster and lightweigh event-driven communication. Exploiting these improved capabilities, a high performance, event-driven archiving system written in C++ has been developed. It inherits the database structure from the existing TANGO Historical Data Base (HDB) and introduces new storage architecture possibilities, better internal diagnostic capabilities and an optimized API. Its design allows storing data into traditional database management systems such as MySQL or into NoSQL database such as Apache Cassandra. The paper describes the software design of the new HDB++ archiving system, the current state of the implementation and gives some performance figures and use cases.

INTRODUCTION

The TANGO archiving system is a tool that allows to store the readings coming from a TANGO based control system into a database. The archived data are essential for the day by day operation of large facilities, such as long term monitoring of subsystems, statistics, correlation of parameters or comparison of operating setups over time.

To take advantage of the fast and lightweight event-driven communication provided by TANGO release 8 [1] with the adoption of ZeroMQ [2], a novel archiving system for the TANGO Controls framework [3] has been designed and developed in collaboration between Elettra and ESRF.

DESIGN GUIDELINES

A number of requirements have been taken into account during the design phase. The HDB++ archiving system must fully comply to the TANGO device server model, with two immediate benefits. First, all the required configuration parameters are stored to and retrieved from the TANGO database; some of these parameters are, for user convenience, duplicated into a dedicated table of the HDB++ schema by a mechanism that guarantees the consistency of the copy. Second, the HDB++ archiving system inherits the TANGO scaling capability: any number of EventSubscriber instances can be deployed according to the desired architecture and overall performance.

The publish/subscribe paradigm is available in TANGO via the event subsystem. More in detail, the *archive* event is provided for archiving purposes and can be triggered on threshold comparison and/or periodic basis. The HDB++ architecture is fully event based; therefore, a part of HDB++ setup consists of conveniently configure TANGO device servers to send events as required.

Two TANGO device servers have been developed. The EventSubscriber, also referenced as archiver, is in charge of gathering the values from the TANGO devices and storing them into the historical database. To address the requirements coming from large systems the need to distribute the workload over a number of archivers shows up. A ConfigurationManager TANGO device server will assist in the operations of adding, editing, moving and deleting an Attribute to/from the HDB++ archiving system. A specific library, exposing a suitable API, addresses the historical data extraction from the archive.

The task of each HDB++ archiving system component is sharply defined; low layer devices, e.g. archivers, have no dependency against the ConfigurationManager, the extraction library or the graphical user interfaces and, possibly, can be deployed standalone. Also, the HDB++ architecture has been designed to easily support different SQL and NoSQL database engines: an abstraction library decouples the interface to the database back-end from the implementation. Adding a new back-end is just matter of writing the code for the specific implementation; this has been done, as an example, during last year to introduce the support for Cassandra [4].

EVENT SUBSCRIBER

The EventSubscriber TANGO device server is the core of the HDB++ archiving system. It subscribes to archive events for the specified Attributes list, stored into a Property in the TANGO database, as well as a number of additional parameters, such as the hostname and port number where the back-end is running, the name of the database and the username and password to be used.

A dedicated thread is in charge of event subscription and callbacks execution; the callbacks, acting as producers, put the complete data structure of the received events in a FIFO queue, protected by a suitable locking mechanism. The thread and the callbacks must be able to handle an arbitrary number of events, possibly limited just by the available memory and the required performances. One additional thread, acting as consumer of the FIFO, is in charge of writing the data into the database. Moreover, a high-mark threshold is setup on the FIFO queue to alert for an overloaded EventSubscriber.

The EventSubscriber device server allows to perform the following operations:

- add/remove an Attribute to/from archiving
- start/stop the archiving for all Attributes
- start/stop the archiving for one Attribute
- read the status of an Attribute
- read the list of Attributes currently archived (started)

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

BIG DATA ANALYSIS & ANALYTICS WITH MATLAB®

D. Willingham#, MathWorks, Sydney, Australia

Abstract

In today's world, there is an abundance of data being generated from many different sources in various industries across engineering, science & business. For example, DESY's 1.7 mile-long PETRA III accelerator is capable of generating 20 gigabyte's of data per second [1]. Using Data Analytics to turn large volumes of complex data into actionable information can help improve design and decision making processes. Big data sets may often be too large for the available memory, take too long to import and process, or just stream too quickly to store. Standard algorithms are usually not designed to process big data sets in reasonable amounts of time or memory usage. MATLAB® is the high-level language and interactive environment used by millions of engineers and scientists worldwide. It lets you explore and visualize ideas and collaborate across disciplines including signal and image processing, communications, control systems, and computational finance. Using a generic case study we demonstrate efficient ways to manipulate, compute and visualise on large, multidimensional data sets on light weight machines. The techniques presented here can also be used to develop predictive models, scale up for high performance computing on clusters, or in the cloud and deployable within databases, Hadoop and other Big Data environments.

INTRODUCTION

Big data describes a massive volume of data that is so large or complex that it's difficult to manage and process using traditional database and software techniques [2]. Big Data is commonly described by the 3 V's

- Volume – which references the amount of data.
- Velocity – is the speed at which data is generated or needs to be analyzed.
- Variety – which describes the range of data types and sources from which data is obtained.

There are a number of challenges that arise when analyzing big data sets, especially for domain experts who are not necessarily experienced software programmers. These include:

- How to get started and gain insight into the basic structure and format of the data especially when the volume exceeds memory limits?
- Rapid exploration of the data, and development of algorithms which can easily be scaled for use with big data
- Using big data algorithms within business systems

As an interesting example of the volume and velocity of data that is being generated in the world. In 2000, the world generated 2.5 Exabyte's of data a day [3]

In this paper, we present a generic case study identifying and predicting patterns of the domestic flights in the US between 1987 and 2008, which is 12gB in size [4]. We address the challenges mentioned previously by demonstrating two capabilities of MATLAB®, namely, datastore and mapreduce and the links with Big Data storage platforms such as Hadoop.

This paper is organised as follows. We highlight 3 techniques that can be used to handle importing and processing big data;

1. Scale the problem down
2. Parallelize the Problem
3. Scale up if needed along with the MATLAB functions that support these.

TECHNIQUES FOR HANDLING BIG DATA

Scale the Problem Down

For simpler problems where data can be too big to be analyzed at once, but where distinct sub-problems can be calculated and the results easily aggregated together, we can simply loop through the various files or records, load the subset of data we're interested in, analyze it, store the result and repeat. MATLAB's `datastore` enables this to be achieved in the easy efficient way.

What is *datastore*?

A *datastore*[4] is an object for reading a single file or a collection of files or data. The *datastore* acts as a repository for data that has the same structure and formatting. For example, each file in a *datastore* must contain data of the same type (such as numeric or text) appearing in the same order, and separated by the same delimiter. A *datastore* is useful when:

- Each file in the collection might be too large to fit in memory. A *datastore* allows the ability to read and analyze data from each file in smaller portions that do fit in memory
- Files in the collection have arbitrary names. A *datastore* acts as a repository for files in one or more folders. The files are not required to have sequential names.
- Not all the data needs to be loaded into memory. For example, you may only wish to analyse 10 columns out of 50 from a data file, so only import the 10 that are needed.

david.willingham@mathworks.com.au

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

TIMING SYSTEMS FOR ATNF TELESCOPES

S. A. Hoyle, CASS, Parkes, NSW, Australia
P. Mirtschin, CSIRO ATNF, Narrabri, NSW, Australia

Abstract

Radio Telescopes require precise time and timing signals for accurate telescope pointing, synchronisation of signal processing instrumentation and offline manipulation of observation data. We provide an overview of the timing system in use at our observatories; briefly describing the main features of the hardware, firmware and software.

INTRODUCTION

The Australia Telescope National Facility, (ATNF) comprises the observatories located near Parkes and Narrabri in NSW and the Murchison Radio Observatory (MRO) in remote Western Australia. The Commonwealth Scientific Industrial Research Organisation (CSIRO) division of Astronomy and Space Science (CASS) manages and operates the ATNF primarily for astronomical research conducted by scientists from institutions in Australia and around the world.

Our radio telescopes vary in age and architecture: from the Parkes 64m single dish (1961), to the Australia Telescope Compact Array (ATCA) (1988) near Narrabri NSW, an array of 6 x 22m diameter moveable antennas, to the Australian Square Kilometre Array Pathfinder (ASKAP) at the MRO, an array of 36 x 12m antennas currently in the outfitting and early commissioning stage of construction. While ASKAP employs the latest technology, our older telescopes have been kept at the forefront of astronomical research due to continual upgrades to their instrumentation and extension of their operating capabilities.

Observations are made with the Compact Array and Parkes 64m telescopes both separately and together to form the Long Baseline Array. Occasionally, the array is extended further with radio telescopes in Tasmania and New Zealand. The technique of using widely separated telescopes in concert provides increased resolution and hence greater detail in the final astronomical images and is known as Very Long Baseline Interferometry (VLBI). Precise time and timing signals are required to synchronise instruments within a single antenna, between antenna elements of an array and between widely separated sites.

Figure 1 illustrates the basic instrumentation of a radio telescope and the points where time and timing signals are applied. Referring to Fig. 1: extremely weak radio frequency energy from a celestial source is converted to an electrical signal and amplified by a multi-band receiver mounted at the focal point of the antenna. The signal is then frequency down-converted and band-pass filtered before being sampled and digitised and passed to the correlator for correlation with itself, and, in the case of an array, with the signals from other antennas. The frequency conversion system requires local oscillators to 'tune' the receiver to the chosen centre frequency within each band and usually other fixed local oscillators for further stages of down conversion to an Intermediate Frequency (IF) suitable for digital sampling. All local oscillators must be precisely controlled to maintain constant phase and phase coherence between the converted signals. Lack of phase coherence results in a reduced output level at the correlator and therefore reduced sensitivity. Phase jitter reduces the quality of the final astronomical image [1]. To this end, the antenna-based oscillators are phase-locked to a highly stable central reference oscillator. In the case of an antenna array, small variations in reference phase at an antenna can occur due to movement and temperature changes in the distribution cabling. These are compensated by measuring the round-trip phase from a central point to each antenna and back; applying corrections either to the correlator output data or to the phase of the local oscillator at the antenna [1].

To attain sufficient dynamic range, the correlator accumulates the correlated input data streams for an 'integration cycle' period, of typically 5 or 10 seconds, prior to transformation to the frequency domain [1]. The integration cycle waveform, along with other real-time sequencing events supplied to the correlator, is generated by an ATNF 'Event Generator'. A Master Clock supplies precise current time to the Event Generator which is pre-programmed to set or clear its digital output event lines at absolute times corresponding to waveform transition edges.

The Antenna Control Computer also uses an Event Generator to generate the 'Start of Integration' events for the Digitiser and the 'Calibration Cycle' waveforms for the Receiver and Conversion System. Start of Integration is encoded in the IF data streams produced by the Digitiser and used to implement variable delays between the signals from an antenna array. The delays are added to compensate for the different arrival times of a celestial signal at each antenna and varied as the earth rotates to ensure alignment of samples from the same wave-front [1]. This concept is illustrated in Fig. 2 below.

The Calibration Cycle is used to determine the overall system gain and hence the true flux of the celestial source. (The system gain itself varies automatically to provide a constant input level to the samplers.) The Calibration Cycle waveform modulates a noise source of known power amplitude at the input to the receiver. The Conversion System then measures the noise on / off power levels and the difference is used to determine the gain [1].

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

VIRTUALISATION WITHIN THE CONTROL SYSTEM ENVIRONMENT AT THE AUSTRALIAN SYNCHROTRON

Ulrich Felzmann, Andrew Starritt, Nicholas Hobbs,
Synchrotron Light Source Australia, 800 Blackburn Road, Clayton 3168, Victoria, Australia

Abstract

Virtualisation technologies significantly improve efficiency and availability of computing services while reducing the total cost of ownership. Real-time computing environments used in distributed control systems require special consideration when it comes to server and application virtualisation. The EPICS environment at the Australian Synchrotron comprises more than 500 interconnected physical devices; their virtualisation holds great potential for reducing risk and maintenance.

INTRODUCTION

The Australian Synchrotron is a 3 GeV third-generation synchrotron light source built in Clayton near the Monash University Campus in Melbourne, Victoria and opened on 31 July 2007. Since 2012, the Australian Synchrotron is operated by the Australian Nuclear Science and Technology Organisation (ANSTO) under the name Synchrotron Light Source Australia (SLSA).

The core of the control system for the 100 MeV linac, the 100 MeV to 3 GeV booster, the 3 GeV storage ring and the 9 beamlines consists of more than 500 Input Output Controllers (IOCs) running EPICS. Most of those IOCs are Linux hosts (mostly CentOS), but also ~ 100 Libera systems, some Windows hosts, VME-bus computers and PLC-based controllers with EPICS embedded.

Each beamline reside in a separate subnet and VLAN and uses its dedicated set of isolated services with one service per host. Due to that philosophy, many of the IOCs do not have any special hardware attached. Those so-called "Soft IOCs" and IOCs with only serial devices attached, that can be easily connected via a Ethernet to serial converter, can be easily converted into virtual machines (VM).

VIRTUAL ENVIRONMENT

Virtualisation, which refers to the creation of an environment that acts like a real computer with an operating system, offers a great way to increase the efficiency and availability of computing services while reducing risk and the total cost of ownership..

In June 2012, the Australian Synchrotron purchased its original hardware for the virtualisation platform. Since then, the initial hardware configuration has been extended substantially and consists currently of the following individual components:

- 4 x Dell R620 hypervisors, providing in total 8 Intel Xeon CPUs (E5-2650 @ 2 GHz) and 768 GB RAM
- 3 x clustered Dell EqualLogic iSCSI devices (104 TB storage group in total), hosting the virtual disk images and frequently changing data for corporate applications that require daily replication (backup)
- 5 x EMC Isilon NL400 nodes (674 TB in total), hosting the virtual disk images for IOCs via NFS as well all the user data from the beamlines
- 2 x CISCO Nexus 5596 switches, connecting the beamlines, corporate and storage networks
- 2 x Juniper EX4300 switches, connecting the machine
- 1 x CISCO Catalyst 3750 switch, connecting the DMZ
- VMware ESXi 6 with resource pooling enabled so "Infrastructure as a service (IaaS)" can be offered to the individual beamlines.

The main components of the virtualisation environment are depicted in Figure **1**.

Apart from services for the control system environment, the virtualisation platform also hosts database and application servers for the corporate business units (email, file server, www, etc.) as well as network management (DHCP, DNS, NTP, PXE, etc.) and monitoring (Observium, Zabbix) services for the facility.

APPLICATIONS

Currently more than 200 virtual machines are active, of which approximately 50 are virtual IOCs (vIOC). In addition, there are 19 EPICS gateways, 3 EPICS CA Archivers (Figure **2**), ~ 15 build boxes, as well as VMs hosting version control, bug tracking and continuous integration testing systems.

Many of the vIOCs have been created from physical machines using the vCenter Converter from VMware, which works successfully most of the time. Some physical hosts with an old Linux operating systems such as. CentOS4 or Debian 3 required a manual conversion.

The median run-time configuration parameters for a vIOC are: 1 vCPU (9.5 MHz), 672 MB RAM allocated, 13 GB disk space used (20 GB thin-provisioned).

Further to hardware virtualisation, some VMs provide further nested virtualisation in the form of operating-system-level virtualisation**.**

- 19 EPICS Gateways providing Channel Access between the different subnets are operated as LXC instances on a single VM to reduce resource consumption.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

STATUS MONITORING OF THE EPICS CONTROL SYSTEM AT THE CANADIAN LIGHT SOURCE

Glen Wright, Michael Bree, Canadian Light Source, Saskatoon, Saskatchewan, Canada

Abstract

The Canadian Light Source (CLS) uses the EPICS Distributed Control System (DCS) for control and feedback of a linear accelerator, booster ring, electron storage ring, and numerous x-ray beamlines. The number of host computers running EPICS IOC applications has grown to over 200, and the number of IOC applications exceeds 700. The first part of this paper will present the challenges and current efforts to monitor and report the status of the control system itself by monitoring the EPICS network traffic. This approach does not require any configuration or application modification to report the currently active applications, and then provide notification of any changes. The second part will cover the plans to use the information collected dynamically to improve upon the information gathered by process variable crawlers for an IRMIS database, with the goal to eventually replace the process variable crawlers.

BACKGROUND

The Canadian Light Source is Canada's national centre for synchrotron research and a global centre of excellence in synchrotron science and its applications. Located at the University of Saskatchewan in Saskatoon, the CLS is a world-class, state-of-the-art facility that is advancing Canadian science, enhancing the competitiveness of Canadian industry and contributing to the quality of life of people around the world.

EPICS is used at the CLS to monitor and control the electronic devices for both electron beam components and photon beam components. While there are many tools that monitor individual Process Variables (PVs) for conditions that indicate problems with individual devices, there are no tools that give an overall picture of the health of the control system itself. Early attempts at the CLS to reflect the status of the control system using an alarm handler meant that the control system was attempting to report on its own health – not an easy feat to accomplish.

Since the CLS first began operation, the monitoring of the control system health was provided by the machine operators in the control room. This was not done as a specific task, but rather occurred when a particular system was unavailable or unresponsive. The natural challenge was not all systems are being used by the operators at all times. A system failure could go unnoticed if it was not in use or directly connected to the PV alarm system. There was a strong need to improve the control system monitoring.

AVAILABLE DATA

It is important to emphasize that the information being collected does not depend on a database of running applications or available PVs. Rather, the purpose of this exercise is to provide maximum status information without having to preconfigure anything.

EPICS uses broadcast packets to indicate the continued running of server processes and when establishing communication between processes. Although the CLS uses individual Virtual Local Area Networks (VLANs) to partition the network, there are currently two systems used for EPICS gateway communication that use 802.1Q packets directly and establish numerous virtual Network Interface Cards (NICs). These two systems can then see all the broadcast packets off of all the configured networks.

The first of the two types of broadcast packets provides the following information:

- Server IP
- Server unique port number
- Packet sequence number

This information is collected to provide a summary as follows:

- The identity of each host running an EPICS server
- The number of EPICS server applications running on any host
- An estimate of how long each server application has been running
- Notification when the number of unique server applications changes
- The timestamp of when the latest packet from a server application had been processed

These packets are sent at regular intervals. These intervals can be configured differently at different sites, but the default is 30 second intervals (more frequent on start-up). Because they are UDP packets, there is a risk that the operating system will drop a packet on a congested network. A single missed packet can't be taken as a problem, but continued missed packets are an alert that something could be wrong.

The second type of broadcast packets come from EPICS client applications, and provides the following information:

- Client IP
- Client Unique TCP port number
- Requested PV name

Collecting this information provides the following summaries:

- Number of times a PV has been requested
- List of clients requesting a PV
- All PV names that have been requested

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

LabVIEW INTERFACE FOR MADOCA II WITH KEY-VALUE STORES IN MESSAGES

T. Matsumoto#, Y. Furukawa, Y. Hamada, T. Matsushita, JASRI/SPring-8, Hyogo, 679-5198, Japan

Abstract

As the next generation of the Message and Database-oriented Control Architecture (MADOCA), MADOCA II is a message-driven distributed control framework that is more accommodating, e.g., it allows control on Windows operating system as well as messages consisting of data of variable [1]. A prototype of MADOCA II was developed in the Laboratory Virtual Instrument Engineering Workbench (LabVIEW) interface in 2013 and implemented on a control system in SPring-8 [2]. However, it was recognized that the interface should be easy-to-use in order to be able to implement MADOCA II in a wide variety of LabVIEW applications, especially in synchrotron radiation experiments. In this paper, we propose a redesigned LabVIEW interface to respond to this challenge. In this interface, messaging is managed through a unified method with virtual instruments (VIs) using key-value stores. As a result, applications for the client program and the equipment management server can be easily constructed. The VIs are based on a dynamic link library (DLL) developed in C++. The interface is easily upgraded by replacing the DLL, which was also applied to the Python interface.

INTRODUCTION

The Message and Database-oriented Control Architecture (MADOCA)-II is the next generation of MADOCA. It was successfully implemented in accelerator controls in the SPring-8 and the SPring-8 Angstrom Compact Free Electron Laser (SACLA) data acquisition (DAQ) system in 2013, as reported at the last International Conference on Accelerator and Large Experimental Physics Control Systems (ICALEPCS) [1]. Like its predecessor [3], MADOCA II is based on a message-oriented control scheme. A message is composed of a text in subject/verb/object/complement (S/V/O/C) syntax. In this context, “S” denotes the program that is automatically defined by the framework and “V” denotes the action of a command, for which “put” or “get” is primarily used. “O” is an object name, which identifies the target of the message, and “C” is an action parameter. Figure 1 shows the structure of MADOCA II messaging. A graphical user interface (GUI) sends a message with “put/camera/init” as V/O/C (S is abbreviated as described), and an equipment manager (EM) of device control responds to the message in the S/V/O/C format. If the initialization of the camera is successful, C is set to “ok” in the response message. Thus, this messaging scheme is easy to understand.

Moreover, MADOCA II is more flexible than MADOCA. In MADOCA II, data strings of varying length, such image data, can be attached to a message, and the Windows environment can be used for device control. These features were utilized in control applications in SPring-8.

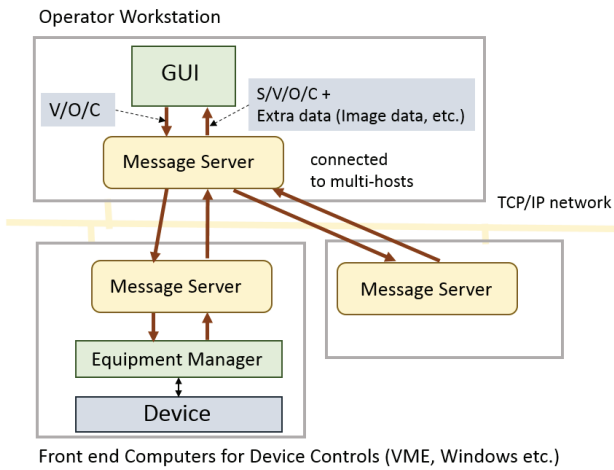


Figure 1: The structure of MADOCA II messaging. Messages are exchanged between a GUI and an equipment manager (EM) thorough a Message Server (MS) in each host.

A motivation underlying the development of MADOCA II was the ability to use the Laboratory Virtual Instrument Engineering Workbench (LabVIEW) interface by using the above flexibilities because LabVIEW is used for some parts in the accelerator control and many applications in synchrotron radiation experiments. If these LabVIEW applications are updated to contain an interface for MADOCA II, we can unify its control procedures in order to save on management costs.

In 2013, we developed a prototype of the LabVIEW interface for MADOCA II, and applied it to a beam position monitor (BPM) with NI PXI-5922 digitizers [2]. In the application, waveform data was attached to a message. A total of 5,000 samples of the waveform could be exchanged between a client application and an EM of the BPM within one second, and the BPM application functioned stably.

We intend to apply the MADOCA II LabVIEW interface to other LabVIEW applications. However, we realize that the interface should be revised for ease of use and better maintainability.

#matumot@spring8.or.jp

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

CUSTOM HARDWARE PLATFORM BASED ON INTEL EDISON MODULE*

D. Pedretti, D. Bortolato, F. Gelain, M. Giacchini, D. Marcato, M. Montis,
S. Pavinato, J. A. Vasquez, INFN-LNL, Legnaro, Italy
M. Bellato, R. Isocrate, INFN Sez. di Padova, Padova, Italy

Abstract

The Computer-on-Module (COM) approach makes cutting edge technology easily accessible and lowers the entry barriers to anyone prototyping and developing embedded systems. Furthermore, it is possible to add all the system specific functionalities to the generic PC functions which are readily available in an off-the-shelf core module, reducing the time to market and enhancing the creativity of system engineers. The purpose of this paper is to show a custom hardware platform based on the tiny and low power Intel Edison compute module [1], which uses a 22 nm Intel processing core and contains connectivity elements to ensure device-to-device and device-to-cloud connectivity. The Intel Edison carrier board designed is expected to act as a local intelligent node, a readily available custom EPICS IOC [2] for extending the control reach to small appliances in the context of the SPES [1] project. The board acts as an Ethernet to RS232/RS422 interface translator with Power-Over-Ethernet supply and network booting as key features of this platform. The x86 architecture of the Edison makes standard Linux software deployment straightforward. Currently the board is in advanced prototyping stage.

INTRODUCTION

Nowadays embedded systems are commonly found in consumer, industrial, automotive, medical, commercial and military applications. An embedded system is hardware and software designed to perform specific tasks but we introduce here a custom hardware platform that is based on a commodity computing platform, thus partially bridging the gap between custom development and commercial off-the-shelf personal computers. The possibility to build an embedded system around a Computer on Module, for example a commercial and highly compact computing platform, represents a step towards a general purpose embedded system.

BOARD OVERVIEW

Particle accelerators need a distributed control network capable to reach devices of different types and requirements. In this context a certain effort is required for embedding the control of a single appliance or a small group of appliances. The custom hardware platform object of this paper has been inspired by the necessity to extend the control reach to small groups of magnet power supplies around the SPES accelerator, actually in construction at LNL (Laboratori Nazionali Legnaro) [3]. Magnet power supply boxes are accessible via serial buses, so we need a low power and low cost microprocessor board running the EPICS IOC software together with one or more RS232/RS422 interfaces. Although the magnet control is the target application, we realized that a more general approach opens the possibility of embedding different instrumentation around the accelerator, preserving the investment.

A Desktop PC

This embedded system is built around the Intel Edison compute module which integrates a 22 nm Intel Atom Processor dual-core 500 MHz. Lot of general purpose PC functionalities like a Wi-Fi dual-band (IEEE 802.11a/b/g/n) controller, an USB 2.0 OTG (On The Go) controller and 40 GPIO (General Purpose Input Output) configurable as I2C, I2S, SPI, SD card, UART, PWM are readily available via a 70 pin board to board connector. This low power Computer on Module is easily found on the market and we decided to use it as the core of our custom platform.

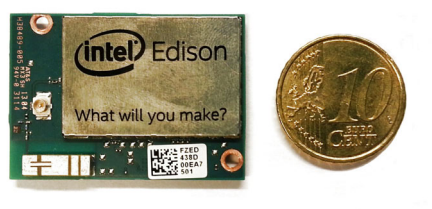


Figure 1: Intel Edison.

Figure 1 shows the tiny processor. Its x86 architecture brings lot of benefits to the embedded controller, which becomes in all respects a custom, low cost, low power and easy to use PC.

Main Features

Figure 2 shows a block diagram resuming the key features of this custom hardware platform. The SPES control system is an Ethernet based distributed network and implements the client-server model foreseen by EPICS architecture. Hence Ethernet connectivity is a key feature of the board. Since Edison Module does not integrate an Ethernet physical layer and neither an Ethernet data link layer we added an Hi-Speed USB 2.0 to 10/100/1000 Gigabit Ethernet controller

* Thanks to Prime Elettronica for assembly the prototype and to DPE Elettronica and Link Engineering for supporting the Place&Route.

[1] SPES (Selective Production of Exotic Species) is a second generation ISOL radioactive ion beam facility. The aim of SPES is to provide high intensity and high-quality beams of neutron-rich nuclei to perform forefront research in nuclear structure, reaction dynamics and inter-disciplinary fields.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

A MULTI-MODAL HUMAN-MACHINE-INTERFACE FOR ACCELERATOR OPERATION AND MAINTENANCE APPLICATIONS

R. Bacher, DESY, Hamburg, Germany

Abstract

The advent of advanced mobile, gaming and augmented reality devices provides users with novel interaction modalities. Today's accelerator control applications do not provide features such as speech, finger- and hand-gesture recognition or even gaze detection. Their look-and-feel and handling are typically optimized for mouse-based interactions and are not well suited for the specific requirements of more complex interaction modalities. This paper describes the conceptual design and implementation of an intuitive single-user, multi-modal human-machine interface for accelerator operation and maintenance applications. The interface seamlessly combines standard actions (mouse), actions associated with 2D single- / multi-finger gestures (touch sensitive display) and 3D single- / multi-finger and hand gestures (motion controller), and spoken commands (speech recognition system). It will become an integral part of the web-based, platform-neutral Web2cToGo framework belonging to the Web2cToolkit suite and will be applicable for desktop and notebook computers, tablet computers and smartphones, and even see-through augmented reality glasses.

INTRODUCTION

In a control room the mouse is the standard user input device to interact with accelerator control applications. Being well accepted by the operators it provides a very accurate pointing capability even onto complex applications with a wealth of graphical widgets. However, hardware commissioning and maintenance use cases might profit from novel interaction capabilities (modalities). For instance the alignment of mirrors mounted on an optical table to adjust a laser beam spot often requires a "third hand". Interacting with control applications via spoken commands could be an appropriate alternative. Likewise remote-controlling a head-mounted display showing some documentation or control application panels might be considerably simplified by recognizing spatial gestures such as clenching a fist or snapping fingers.

This paper reports the conceptual design and the implementation of a single-user human-machine-interface (HMI) for accelerator operation and maintenance applications in the context of the Web2cToolkit Web service collection [1,6]. Web2cHMI defines a set of common user interactions associated with various modalities including flat or spatial gestures, spoken commands, and ordinary mouse or touch actions. This approach applies to desktop or notebook computers with passive or touch-sensitive display, tablet computers or smartphones, and even see-through Augmented Reality glasses.

Web2cToolkit is a collection of Web services including among others Web2cViewer and Web2cToGo. Both Web applications implement Web2cHMI. The Web2cViewer Web service provides a user-configurable interactive synoptic live display to visualize and control accelerator or beam line equipment. The Web2cToGo client application is especially designed for mobile devices and is capable of running simultaneously up to 15 multi-page Web2cToolkit-compliant Web applications such as Web2cViewer applications. At any particular time the selected page of the selected user application is displayed while other pages of the same application and the pages of the other applications are hidden. The Web2cToGo client application provides three different views including Explorer View, Navigation View, and Operation View.

CONCEPTUAL DESIGN AND BASIC FEATURES

Web2cHMI is a set of JavaScript (JS) classes to be used in Web client applications executed by a Web browser. A Web2cHMI-compliant browser must implement the Web Sockets API and the getUserMedia / Stream API (required for speech recognition only). Table 1 summarizes the implementation status [2] by major Web browsers (most recent desktop or mobile version).

The recognized user input from each supported modality is handled by its own modality-specific class notifying a common interface class which translates the modality-specific user actions into common unique application-specific commands for further execution.

Table 1: Implementation Status of Required JS API

Web Browser	Web Sockets API	getUserMedia / Stream API
Edge	yes	yes
Firefox	yes	yes
Chrome	yes	yes
Safari	yes	no
Opera	no	yes
iOS Safari	yes	no
Opera Mini	no	no
Android Browser	yes	yes
Chrome for Android	yes	yes

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

A GRAPHICAL TOOL FOR VIEWING AND INTERACTING WITH A CONTROL SYSTEM

J. Forsberg, V. Hardion, D. Spruce
MAX IV Laboratory, Sweden

Abstract

We present a graphical interface for displaying status information and enabling user interaction with the Tango based control system for the MAX IV synchrotron. It focuses on bringing an intuitive view of the whole system, so that operators can quickly access the controls for any hardware based on its physical location.

The view is structured into different layers that can be selectively shown, and various live updated information can be displayed in the form of e.g. color or text. Panning and zooming is supported, as well as invoking commands. The interface is defined by an SVG (Scalable Vector Graphics) drawing which can be edited without programming expertise.

Since our system is based on modern web technologies, it can be run as a web service accessible by standard browsers, but it can also be integrated in GUI applications.

INTRODUCTION

The MAX IV laboratory is a synchrotron facility under construction in Lund, Sweden. It has a TANGO [1] based control system consisting of thousands of different pieces of software that provide users with access to equipment. The control system tends to be structured in a way that makes sense to the software developers that build it, but not necessarily to accellerator physicists and synchrotron operators. This project, called "SVG synoptic", is an attempt at making a control system interface that is intuitive to the users.

The basic idea is to display live control system information in the form of a picture that somehow represents the physical layout of the machine, in a schematic way. Items in the picture may contain live information about corresponding control system variables, as e.g. color or text. The user is able to iteract with the drawing in order to bring up specific controls in a window, or to perform specific actions.

It may not be realistic to fit every detail of a complex system into a single image. By allowing the synoptic to contain several different detail levels, we can provide a seamless interface where the user can choose what to see simply by zooming and panning using the mouse.

A control system is usually logically split into several separate sub-systems such as vacuum, cooling, and so on, and for many use cases only some of these may be of interest to the user. SVG synoptic provides a way to structure the information into layers that can be toggled on or off depending on what subsystems are relevant at the time.

The size of the control system has to be taken into account from a performance point of view. Communicating with thousands of items is time consuming, and since most of the information is not visible at any given time it is not helpful. Therefore, SVG synoptic has the ability to only keep updated about the parts currently visible to the user.

CHOICE OF TECHNOLOGY

The development of this project was initially motivated by requirements from the users that were not easily achievable using available tools. The choice was between extending an existing system or creating a new one.

SVG

The standard tool for displaying synoptics in the TANGO community is based on a specialized Java based drawing tool called JDraw (part of the ATK [2] suite of tools). It has some of the features of SVG-synoptic, but the current viewers can only provide a static view. The JDraw format, while simple, is also fairly limited. Since it is not a widely used standard, there is not much tooling to work with the file format.

SVG (Scalable Vector Graphics) [3] has the following characteristics:

- Being a vector format (like JDraw) it produces resolution independent images suitable for free zooming.
- It is a mature, widely used and open standard, adopted by web browsers (currently all major browsers except IE 8 or older) and many graphics software packages.
- The file format is XML and therefore straightforward to generate and manipulate programmatically.
- SVG can easily be accessed from javascript since it becomes part of the DOM (Document Object Model) once loaded. It can also be styled through CSS (Cascading Style Sheets).
- The FOSS (Free and Open Source Software) drawing application Inkscape [4] uses SVG as a native format.

Considering these features, we chose to go with SVG instead of extending JDraw.

JavaScript

Part of the reasons SVG was chosen as an image format was that it enables the use of a web browser as viewer, even if embedded in an application. This cut down development time, but it did mean that the user interaction had to be implemented in JavaScript [5], the "native" programming language supported by web browsers.

JavaScript has a large community around it, resulting in a lot of high quality libraries. In particular, for interacting

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

HOW CASSANDRA IMPROVES PERFORMANCES AND AVAILABILITY OF HDB++ TANGO ARCHIVING SYSTEM

R. Bourtembourg, J.L. Pons, P. Verdier, ESRF, Grenoble, France

Abstract

The TANGO release 8 led to several enhancements, including the adoption of the ZeroMQ library for faster and lightweight event-driven communication. Exploiting these improved capabilities, a high performance, event-driven archiving system, named Tango HDB++, has been developed. Its design gives the possibility to store archiving data into Apache Cassandra: a high performance scalable NoSQL distributed database, providing High Availability service and replication, with no single point of failure. HDB++ with Cassandra will open up new perspectives for TANGO in the era of big data and will be the starting point of new big data analytics/data mining applications, breaking the limits of the archiving systems which are based on traditional relational databases. The paper describes the current state of the implementation and our experience with Apache Cassandra in the scope of the Tango HDB++ project. It also gives some performance figures and use cases where using Cassandra with Tango HDB++ is a good fit.

INTRODUCTION

The TANGO archiving system is a tool allowing TANGO users to store the readings coming from a TANGO based control system into a database. The archived data are essential for the day by day operation of complex scientific facilities. They can be used for long term monitoring of subsystems, statistics, parameters correlation or comparison of operating setups over time.

To take advantage of the fast and lightweight event-driven communication provided by TANGO release 8 [1] with the adoption of ZeroMQ [2], a novel archiving system for the TANGO Controls framework [3], named HDB++ [4], has been designed and developed, resulting from a collaboration between Elettra and ESRF.

HDB++ design allows TANGO users to store data with microsecond timestamp resolution into traditional database management systems such as MySQL [5] or into NoSQL databases such as Apache Cassandra [6].

APACHE CASSANDRA

Apache Cassandra is an open source distributed database management system available under the Apache 2.0 license. Cassandra's masterless ring architecture where all nodes play an identical role, is capable of offering true continuous availability with no single point of failure, fast linear scale performance and native multi-datacentre replication. All these advantages over the traditional relational databases made Cassandra a very good candidate to store the historical data coming from the ESRF TANGO control systems.

HDB++ DESIGN

HDB++ is a novel TANGO Controls archiving system. Its architecture is mainly based on the two following Device Servers and takes advantage of the TANGO events mechanism:

- The Event Subscriber Device Server subscribes to TANGO archive events, which are ZeroMQ events in the latest TANGO releases, and stores the received events in the historical database. It provides diagnostics data as well.
- The Configuration Manager Device Server configures the attributes to be archived and defines which Event Subscriber is responsible for a set of TANGO attributes to be archived. It provides diagnostics data as well.

An abstraction library named *libhdb++* decouples the interface to the database back-end from the implementation.

To be able to store data into Apache Cassandra, a C++ library named *libhdb++cassandra*, implementing the methods from the *libhdb++* abstract layer for Cassandra, has been developed at the ESRF. It was inspired by the work done by Elettra who implemented the HDB++ MySQL back-end support shared library. The HDB++ Device Servers design is shown in Fig. 1.

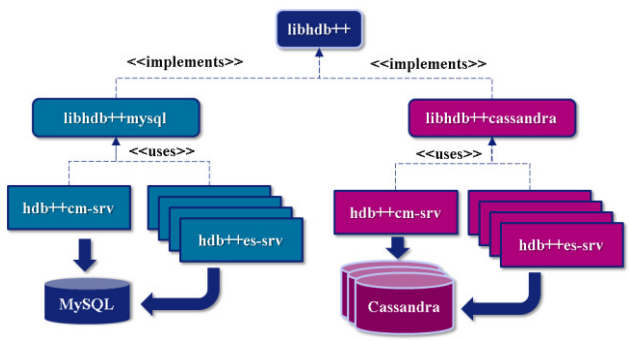


Figure 1: HDB++ Device Servers design.

The Cassandra HDB++ libraries allow us to reuse the Event Subscriber and Configuration Manager Device Servers as well as the Graphical User Interfaces (GUI) interacting with these Device Servers (HDB++ Configurator GUI shown on Fig. 2 and HDB++ Diagnostics GUI), without changing their source code. Linking the two Device Servers with *libhdb++cassandra* was sufficient to be able to store data into Cassandra.

For the data extraction part, new Java classes were developed for HDB++ Cassandra, extending the features provided by an abstract layer hiding the database back-end to the Java applications.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

THE INSTRUMENT CONTROL ELECTRONICS OF THE ESPRESSO SPECTROGRAPH @VLT

V. Baldini, G. Calderone, R. Cirami, I. Coretti, S. Cristiani, P. Di Marcantonio, P. Santin, INAF-OATs, Trieste, Italy
D. Mégevand, OG, Versoix, Switzerland,
F. Zerbi, INAF-OAB, Merate, Italy

Abstract

ESPRESSO, the Echelle SPectrograph for Rocky Exoplanet and Stable Spectroscopic Observations, is a super-stable Optical High Resolution Spectrograph for the Combined Coudé focus of the Very Large Telescope (VLT). It can be operated either as a single telescope instrument or as a multi-telescope facility, by collecting the light of up to four Unit Telescopes (UTs). From the Nasmyth focus of each UT the light is fed, through a set of optical elements (Coudé Train - CT), to the Front End Unit (FEU) which performs several functions, as image and pupil stabilization, inclusion of calibration light and refocusing. The light is then conveyed into the spectrograph fibers.

The whole process is handled by several electronically controlled devices. About 40 motorized stages, more than 90 sensors and several calibration lamps are controlled by the Instrument Control Electronics (ICE) and Software (ICS). The technology employed for the control of the ESPRESSO subsystems is PLC-based, with a distributed layout close to the functions to control. This paper illustrates the current status of the ESPRESSO ICE, showing the control architecture, the electrical cabinet's organization and the experiences gained during the development and assembly phase.

INTRODUCTION

ESPRESSO is the Echelle SPectrograph for Rocky Exoplanet and Stable Spectroscopic Observations that will be installed in the Combined Coudé Laboratory (CCL) of the ESO VLT. The main goals of ESPRESSO will be the measurement of high precision radial velocities of solar type stars for search for rocky planets, the measurement of the variation of the physical constants and the analysis of the chemical composition of stars in the nearby galaxies [1]. To reach these goals an extremely stable, accurate, efficient and high resolution spectrograph is required.

ESPRESSO will be able to operate either as a single telescope instrument or as a multi-telescope facility, collecting the light of up to four telescopes. In fact the light of each UT is fed from the Nasmyth focus, through a tunnel made up of optical elements (Coudé Train), to a Front End Unit (FEU) in the CCL of ESO VLT.

The ESPRESSO Instrument Control Electronics aims to control all moving parts that allow the light to follow that path and the lamps that calibrates the final spectrum, assuring also the safety of the instruments and people through a large number of sensors and alarms.

SYSTEM ARCHITECTURE

Each of the FEU four arms (see Figure 1) is composed by an Atmospheric Dispersion Corrector (ADC), a focus translational stage and a system that performs the field and pupil stabilization. The latter is composed by a set of piezo tip-tilt stages to modify the light path, two technical CCD (TCCD) and a neutral filter. The FEU provides, through a calibration slide, the means to inject the calibration light (white and spectral sources) into the spectrograph fiber when needed.

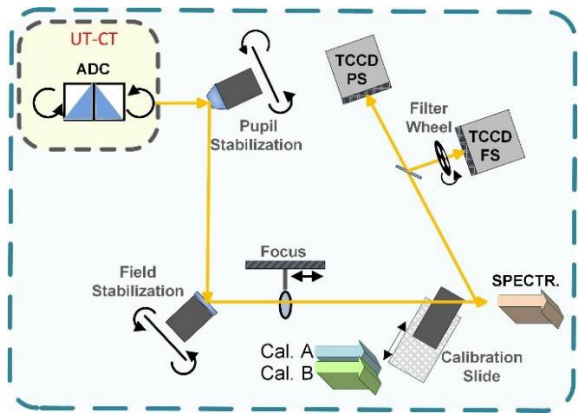


Figure 1: FEU functions in each of the four arms.

Moreover, the FEU provides a mode selector (see Figure 2) mounted on a rotary stage whose task is to feed the spectrograph fibers with the light coming from the selected telescopes and passing through the corresponding FEU arms in one of the specific instrument mode (single telescope High Resolution, single telescope Ultra High Resolution or multi telescope Medium Resolution).

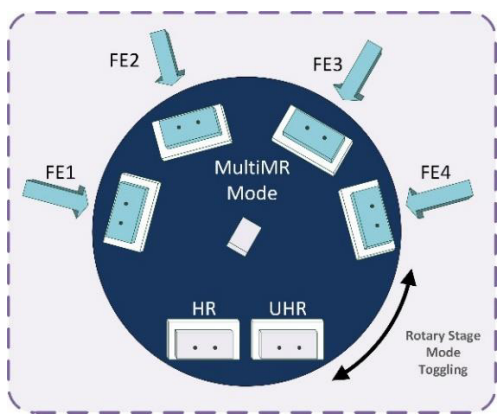


Figure 2: ESPRESSO mode selector.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

A PROTOCOL FOR STREAMING LARGE MESSAGES WITH UDP

C. Briegel, Rich Neswold, Mike Slicznia
FNAL†, Batavia, IL 60510, U.S.A.

Abstract

We have developed a protocol concatenating UDP datagrams to stream large messages. The datagrams can be sized to the maximum of the receiver. The protocol provides acknowledged reception based on a sliding window concept. The implementation provides for up to 10 MByte messages and guarrantees complete delivery or a corresponding error. The protocol is implemented as a standalone messaging between two sockets and also within the context of Fermilab's ACNet protocol. The result of this implementation in vxWorks is analysed.

INTRODUCTION

The Fermilab control system [1, 2] required an alternative to providing large messages. Currently, the user must piece fragments together to support large messages. An atomic message needed to be transmitted guaranteeing delivery and provide the user with an atomic operation.

A working group was formed and discussed several alternatives. The resolution was to support existing protocols but with an extension protocol to concatenate messages with UDP datagrams. The user would not have to modify any code and realize a significantly larger message size. In comparison with other strategies, the implementation provided a solution with little impact to the control system, transparent to the user, minimized risk, and could be accomplished with minimal effort.

First, the implementation was accomplished as a standalone protocol. A simple interface was established to send large messages from one node to another with call-backs when the transmitter or receiver was complete. This minimal test solution would provide a base line for performance as well as a proof-of-principle for the protocol

Also, the protocol was used to increase the message size for ACNet [3], the Fermilab control system's messaging protocol. The added protocol for ACNet was implemented such that all existing protocols could be transparently supported by large messages. This included requests, replies to requests, and unsolicited messages.

PROTOCOL

The protocol in Figure 1 provides a mechanism to concatenate datagrams together. By acknowledging a sliding window of frames received, the message can be delivered reliably to the destination. The protocol abides by the following rules and recommendations:

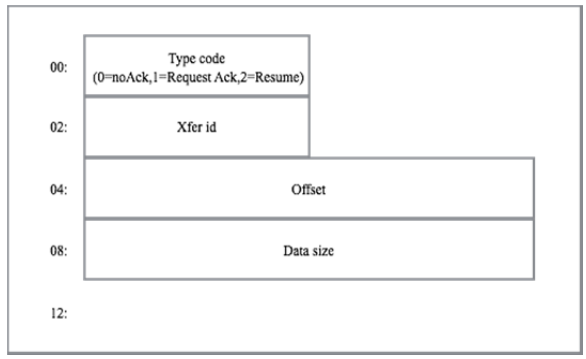


Figure 1: Large Message Protocol.

Rules

- If the offset is zero, then a new reply is arriving. The receiver can use the data size field to pre-allocate a buffer to hold the rest of the incoming data. After saving the data in the buffer, it sets the next expected offset to be equal to the size of data that was just received.
- If the offset is non-zero, it checks to see if the offset and transfer ID matches a reply that is in progress. If a match is found, the data is appended to the buffer and the next expected offset is updated.
- After appending the data, if the packet also asked for a response (type code 1 in the long message header), the task will send a resume message (Figure 2) with the current expected offset.
- If the offset is non-zero and a reply to a transfer ID is in progress but the offset is too high (a packet was dropped), the task waits for a packet that also wants a reply. When it arrives, a resume message is sent to the sender with the offset of the missing data.
- When the transfer is complete, the last packet will also require a response. The receiver returns the expected offset (which at this point will be the size of the data) or a previous offset, if a packet was dropped.

Recommendations

- The first segment **should** use type code 1, asking the receiver for a resume message. By doing this, part of the payload gets sent in addition to checking whether the receiver supports large messages (a timeout indicates no support.)
- The last packet of the message **should** use type code 1 to make sure the entire message was received.
- The sender **may** vary the interval between ACK requests to adapt to network conditions. For instance, the sender might begin the transfer with an interval of 4 packets before asking for an ACK. If there isn't an error, then 8 packets can be sent before the next

†Operated by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the United States Department of Energy.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

THE NEW MODULAR CONTROL SYSTEM FOR POWER CONVERTERS AT CERN

M. Di Cosmo, B. Todd
CERN, Geneva, Switzerland

Abstract

The CERN accelerator complex consists of several generations of particle accelerators, with around 5000 power converters supplying regulated current and voltage to normal and superconducting magnet circuits. Today around 12 generations of converter control platforms can be found in the accelerator complex, ranging in age and technology. The diversity of these platforms has a significant impact on operability, maintenance and support of power converters. Over the past few years a new generation of modular controls called *RegFGC3* has been developed by CERN's power conversion group, with a goal to provide a standardised control platform, supporting a wide variety of converter topologies. The aim of this project is to reduce maintenance costs by decreasing the variety and diversity of control systems whilst simultaneously improving the operability and reliability of power converters and their controls. This paper describes the state of the on-going design and realization of the RegFGC3 platform, focusing on functional requirements using Thyristor based converters as an example.

INTRODUCTION

The CERN accelerator complex consists of several different machines, each with particular powering characteristics. Power converters are a fundamental requirement of any such machine, being used to provide power to various loads such as magnets (normal, superconducting, dipole, quadrupole …), radio-frequency (RF) cavities and other powered equipment. Evolutions in both accelerator functional requirements and power technology require a regular updating of approaches for converters and their associated control electronics, in particular with enhanced diagnostics and more demanding regulation performance.

Table 1 gives a break-down of typical power converter control platforms, and their use in the CERN accelerator complex.

This diversity of platforms gives rise to significant challenges:

- The age of legacy platforms is expected to give rise to end-of-life effects, where some platforms are already >30 years old.
- Information about these platforms is not always coherently stored and available for experts.
- Some platforms cannot be re-produced, due to component obsolence.
- Each control platform generally has an associated software stack, which needs to be maintained.
- The engineers and system responsibles for these systems are not always available to help in diagnosing faults and issues.
- Operations increasingly want better control over some of these platforms, working on them is becoming increasingly difficult and risky.

This leads to higher failure rates, longer repair times, and generally lower availability of the CERN accelerator complex.

Table 1: Control Systems and Converters at CERN

Control Platform	Machine / Area	Number of converters
MIL1553	PS	1218
PLC	SPS, AD, LIER	510
RS422	PS	366
Junction Crate	North Area	322
MUGEF	SPS, TTs	652
FGC2	LHC	1782
RegFGC3	Various	497

A NEW CONTROL SYSTEM

In recent years, new projects and consolidation activities led to the development of a new control platform, as existing solutions were not able to meet requirements. This presented itself as an opportunity to develop a new approach for power converter controls, whereby a single platform could be developed to provide a standard, modular solution for any past, present or future converter at CERN. This new platform is called Regulation based on the 3rd Generation Function Generator Controller (RegFGC3).

The aim of this platform is to provide for the new projects (e.g. the LHC Injector Upgrade and LINAC4) but also to be applied to any requirements in the coming years [1, 2]. The medium term goal is to reduce the number of control systems in order to improve system maintainability and operability, reducing the burden on the controls sections in the power converter group. In this framework, the EPC group developed the RegFGC3.

THE REGFGC3 CONTROL SYSTEM

The RegFGC3 is a modular converter control platform developed at CERN with the main goal of providing a standardised solution and satisfying as many requirements

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

MAGNET SERVER AND CONTROL SYSTEM DATABASE INFRASTRUCTURE FOR THE EUROPEAN XFEL

Lars Fröhlich, Piotr Karol Bartkiewicz, Marcus Walla, DESY, Hamburg, Germany

Abstract

The linear accelerator of the European XFEL will use more than 1400 individually powered electromagnets for beam guidance and focusing. Front-end servers establish the low-level interface to several types of power supplies, and a middle layer server provides control over physical parameters like field or deflection angle in consideration of the hysteresis curve of the magnet. A relational database system with stringent consistency checks is used to store configuration data. The paper focuses on the functionality and architecture of the middle layer server and gives an overview of the database infrastructure.

INTRODUCTION

The European X-ray Free Electron Laser (XFEL) is a research facility currently under construction in close collaboration between the European XFEL Facility GmbH[1] and DESY[2] in Hamburg, Germany [1–3]. It consists of a superconducting linear accelerator delivering an electron beam with particle energies up to 17.5 GeV and a total beam power up to ~600 kW, several long undulator sections enabling the generation of extremely brilliant X-ray pulses at wavelengths down to 0.05 nm, beamlines and setups for photon science experiments, and the associated infrastructure.

The electron beamlines of the European XFEL will be equipped with more than 1400 individually powered electromagnets for beam guidance and focusing. In this paper, we give an overview of the various control system components that work together to provide a uniform interface for controlling and monitoring these magnets. We start at the interface between hard- and software with a short introduction to the server architecture of the magnet power supply system. Above this level, a new magnet middle layer server has been introduced that will be treated in more detail. It gives access to physical parameters like magnetic fields and deflection angles in addition to handling hysteresis effects. The paper closes with a few remarks on the database back-end used for the storage of configuration parameters for these and other systems.

POWER SUPPLY SERVERS

As shown in Fig. 1, distributed front-end servers establish a low-level interface to the magnet power supply controllers (PSCs). There are currently two independent PSC families using different communication protocols and interfacing logic. Most of the main magnets (bending dipoles, quadrupoles) along with some correctors are controlled by PSCs designed at DESY using the industrial CANopen communication protocol [4]. The PSCs of the second family originate from Budker Institute of Nuclear Physics (BINP), Novosibirsk, Russia, and use a proprietary protocol [5] which is also CAN-based.

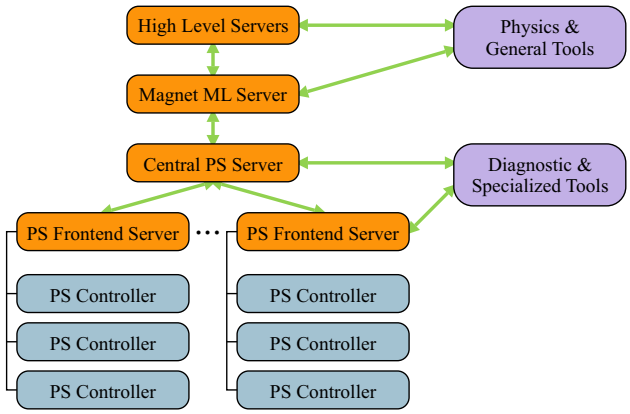


Figure 1: An overview of communication paths between user tools, various server layers, and power supply controller hardware.

An important task of the front-end servers is to hide differences in the interfacing logic and to provide a unified command set that is independent of hardware and communication protocols. The command set covers primitive actions such as switching a power supply on/off, ramping output current, setting current limits, delivering status and error information, and so on. The front-end servers also hide all addressing details, publishing descriptive names of the PSCs rather than their CAN bus identifiers and numbers. The front-end servers can be accessed by a central power supply server (CPSS) as well as by diagnostic applications using the TINE [6, 7] protocol.

The CPSS provides the only access to magnet power supplies for middle layer servers or other clients. It dispatches all requests referring to a single PSC to the corresponding front-end server where they are converted to PSC-specific commands. The CPSS is also capable of accepting composite commands that address PSC groups rather than single units. Such commands are split by the CPSS into several primitive commands and delivered to the front-end servers. The CPSS continuously monitors the front-end servers and collects status information, making it available for middle layer servers and other diagnostic or archiving tools.

MAGNET MIDDLE LAYER SERVER

For previous accelerators at DESY, the main magnet-related quantity provided at the control system level was

[1] European X-Ray Free-Electron Laser Facility GmbH, Albert-Einstein-Ring 19, 22761 Hamburg, Germany
[2] Deutsches Elektronen-Synchrotron DESY, Notkestraße 85, 22607 Hamburg, Germany

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

SECURING ACCESS TO CONTROLS APPLICATIONS WITH APACHE HTTPD PROXY

Piotr Golonka, CERN, Geneva, Switzerland,
Hannu Kamarainen, Jyväskylä University of Applied Sciences, Institute of Information Technology, Jyväskylä, Finland

Abstract

Many commercial systems used for controls nowadays contain embedded web servers. Secure access to these, often essential, facilities is of utmost importance, yet it remains complicated to manage for different reasons (e.g. obtaining and applying patches from vendors, ad-hoc oversimplified implementations of web-servers are prone to remote exploit). In this paper we describe a security-mediating proxy system, which is based on the well-known Apache httpd software. We describe how the use of the proxy made it possible to simplify the infrastructure necessary to start WinCC OA-based supervision applications on operator consoles, providing, at the same time, an improved level of security and traceability. Proper integration with the CERN central user account repository allows the operators to use their personal credentials to access applications, and also allows one to use standard user management tools. In addition, easy-to-memorize URL addresses for access to the applications are provided, and the use of a secure https transport protocol is possible for services that do not support it on their own.

INTRODUCTION

Controls applications for numerous systems at CERN's Large Hadron Collider, as well as for technical infrastructure are built with the WinCC Open Architecture SCADA software [1]. Being highly modular and inherently distributed in nature, it allows the separation of the operator consoles (typically in the Control Rooms, or remote access using Terminal Servers) executing the "User Interface" part of the controls application from other critical tasks performed on the control servers. This architecture allows for good scalability in terms of the number of users. A typical problem is however to deliver files, such as panels, libraries or binary plugins, needed at run time, to the operator consoles. These files are application-specific and therefore need to be stored in the control servers where the controls application is running. With over 500 consoles in the control rooms, 20 Terminal Servers and ~200 controla applications that should all be accessible from any console enforcing access control to protect from unauthorized use becomes complex to handle.

In the classical approach applied until now, a file-share is created on every control server, for every application, to make the necessary server-stored files accessible on every console. The fact that files need to be shared to clients running Linux (for consoles) and Windows (for Terminal Servers) operating systems makes it necessary to maintain two parallel file-sharing technologies (NFS and SAMBA), see Figure 1A. Even with an administrative toolbox, support from the operating system administrator, and automated tools to generate application-specific configuration files for the UIs, the management and maintenance of an ever-growing number of applications becomes a time-consuming task. The hardship of resolving numerous performance and stability problems have also prompted us to search for alternative ways of starting the application UIs on operator consoles.

Recent versions of WinCC OA allow for an alternative way of distributing the files to the User Interfaces using the *http* protocol. The files are served by the embedded http server component of WinCC OA, and may be requested as necessary by the clients, who also maintain a cache of the files on a local disk to improve performance. Optionally, the connection to process-data could also be tunnelled by this server, yet we do not use it. In this setup, any machine (with WinCC OA installed) could request the necessary files and hence run the controls application, by simply knowing the appropriate *URL* (address) at which the control server exposes the necessary files. Even though some basic security mechanisms are built into the WinCC OA web server, they are suitable for isolated plants and would not scale for large scale deployments. Moreover, configuring the TLS certificates to obtain a secure https connection is hard to maintain, as well as the fact that there is typically more than a single controls application running on a server. In addition, in the basic setup, we would not be able to assure proper functioning of redundant systems.

In the remaining part of the paper we will present an architecture with a http proxy server in between the server and the UI part of the application, and we will describe the benefits of such arrangement.

ARCHITECTURE WITH A REVERSE PROXY SERVER

A proxy server is a well-known pattern in computer networks, particularly for the World Wide Web: a dedicated server acts as an intermediary between clients and a server that hosts some services. Of our interest are the reverse proxies, that protect the server from uncontrolled access to its services coming from a large network, and perform tasks such as authentication/authorization, load-balancing or certificate handling (SSL/TLS offloading).

The popular Apache *Httpd* web server [2] readily available on Linux platforms, may be configured to act as

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

PROGRESS OF THE CONTROL SYSTEMS FOR THE ADS INJECTOR II*

Yuhui Guo, Haitao Liu, Jing Wang, Jiangbo Luo, Yongpeng Wang, Zhiyong He,Ting Liu
Institute of Modern Physics, Chinese Academy of Sciences, Lanzhou, China

Abstract

This paper reports the progress of the control system for accelerator injector II used in China initiative accelerator driven sub-critical (ADS) facility. As a linear proton accelerator, injector II includes an ECR ion source, a low-energy beam transport line, a radio frequency quadrupole accelerator, a medium energy beam transport line, several crymodules, and a diagnostics plate. Several subsystems in the control system have been discussed, such as a machine protection system, a timing system, and a data storage system. A three-layer control system has been developed for injector II. In the equipment layer, the low-level control with various industrial control cards, such as programmable logic controller and peripheral component interconnect (PCI), have been reported. In the middle layer, a redundant Gigabit Ethernet based on the Ethernet ring protection protocol has been used in the control network for injector II. In the operation layer, high-level application software has been developed for the beam commissioning and the operation of the accelerator. Finally, by using this control system, the proton beam commissioning for injector II in the control room has been mentioned.

INTRODUCTION

The Chinese Academy of Sciences initiated an accelerator driven sub-critical (ADS) program in 2011 under the frame of "Strategic Priority Research Program" for the objective of the safe disposal of nuclear waste as well as the potentials for advanced power generation [1]. In an ADS system (e.g. [2]-[3]), the high intensity beams from a proton accelerator bombard a heavy-metal spallation target located at the centre of a sub-critical reactor core to produce an intense neutron source. Then, the sub-critical reactor is driven by this intense external neutron source from the spallation target to achieve criticality.

In the China initiative ADS system, the proton accelerator consists of two injectors, each for the energy of 25 MeV, and a main accelerator which is designed for the energy of 250 MeV and the current of 10 mA. A demo facility for injector II has been manufactured in the Institute of Modern Physics of Chinese Academy of Science. In this paper, we report the control system for the demo facility of injector II.

* This work is supported by Strategic Priority Research Program of Chinese Academy of Sciences (XDA03021503)
Corresponding author, Email: guoyuhui@impcas.ac.cn

THE OVERALL ARCHITECTURE OF CONTROL SYSTEM

Injector II in the China ADS system includes an ECR ion source, a low-energy beam transport line (LEBT), a radio frequency quadrupole accelerator (RFQ), a medium energy beam transport line (MEBT), 4 crymodules and a diagnostics plate (Dplate). Each crymodule consists of 6 superconducting Half Wave Resonators (HWR) cavities, 6 superconducting solenoids, and 5 beam position monitors (BPM). The main responsibility of the control system for injector II is to operate the accelerator injector II in the normal and harmonious fashion, and to monitor various parameters during the operation of the accelerator. The stability, reliability and ease of use are the main issues while designing the control system for injector II. Based on the Epics system architecture, a general three-layer control system is constructed for the ADS injector II [4].

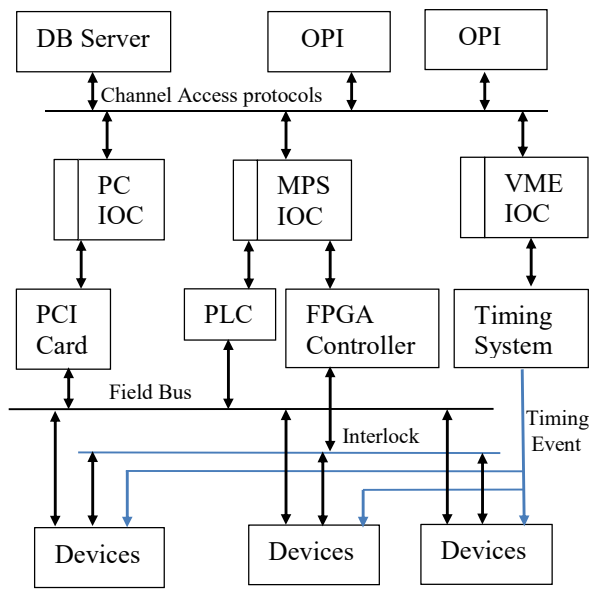


Figure 1: Overall structure of control system for injector II.

As shown in Fig. 1, at the top layer there is the OPI and DB servicer layer. It accesses two low-level machines for interfacing Epics IOC and the field devices by the intranet. At the same time, machine protection system and timing system will be built on the network of real-time control system with optical fiber. Based on the technology of FPGA and PCI bus, fast interlock controller can protect the critical equipment components from damage. The remote control and parameter setting for the field devices is implemented by the PCI cards and the PLC system.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

INFORMATION SECURITY ASSESSMENT OF CERN ACCESS AND SAFETY SYSTEMS

T. Hakulinen, X.B. Costa Lopez, P. Ninin, P. Oser
CERN, Geneva, Switzerland

Abstract

Access and safety systems are traditionally considered critical in organizations and they are therefore usually well isolated from the rest of the network. However, recent years have seen a number of cases, where such systems have been compromised even when in principle well protected. The tendency has also been to increase information exchange between these systems and the rest of the world to facilitate operation and maintenance, which further serves to make these systems vulnerable. In order to gain insight on the overall level of information security of CERN access and safety systems, a security assessment was carried out. This process consisted not only of a logical evaluation of the architecture and implementation, but also of active probing for various types of vulnerabilities on test bench installations.

INTRODUCTION

Information Security of Control Systems

While network accessible server and desktop systems are nowadays out of necessity fairly well secured, control systems have until recent times been largely ignored when it comes to information security. This is due to the fact that most such systems are by their nature private and disconnected. Control system vendors also tend to strongly recommend that their systems be kept in isolation of publicly available networks, for which reason neither the vendors nor the clients may have seen the need for rigorous security. However, tendency nowadays is to increase information exchange between control systems and the rest of the organization to allow, e.g., supervision and remote control, and to facilitate operation and maintenance. Recent events, such as the infamous Stuxnet episode [1], have also served to demonstrate that even well isolated systems are not necessarily immune to security problems.

Access and Safety Systems

CERN GS/ASE group is responsible for the management of CERN personnel access and safety systems. This includes specification, design, implementation, contracting, operation, and maintenance of these systems. The current principal access and safety systems under our responsibility are LHC Access Control System (LACS), LHC Access Safety System (LASS), PS Access Control System (PACS), PS Access Safety System (PASS), Personnel Safety System of the SPS complex, Surveillance of Sites system (SUSI), CERN Safety Alarm Monitoring system (CSAM), Gas detection and alarm system (Sniffer), Site Information Panels / Simple Access Messages system (SIP/SAM), and Safety System Atlas (SSA).

CERN access and safety systems typically consist of many different kinds of devices, such as Windows and Linux servers, operator posts, panel-PCs, PLCs, video cameras, interphones, card readers, biometry scanners, etc. These devices come from many different vendors, and they are nowadays mostly directly network connected. Access systems reside mainly in the restricted CERN Technical Network (TN) but some devices reside also in the CERN-public General Purpose Network (GPN). The most important systems have their own private networks. Most systems also have their respective test platforms, which aim to replicate the production systems in sufficient detail, albeit in a smaller scale, in order for development and testing to be possible without endangering production.

We carried out an information security assessment of two of our most visible access systems, LACS and PACS. The work consisted first of creating an inventory of all the devices and services, which included establishing dependencies between systems, and then carrying out active penetration testing of the various services using available security tools or writing specific utilities, whenever necessary. The goal was not only to find any existing vulnerabilities, but also to gain a larger understanding of the principal risks involved and to develop procedures for managing such risks.

INVENTORY

The first task in a security assessment of a complex system is to create a comprehensive inventory of the entire system: types of devices, operating systems they run, services they use and provide, interdependencies they have. After all the devices were identified and categorized, representatives of the most interesting devices within a category were selected and focused on. Fifteen generic groups were identified within the test systems: servers, access point PCs, operator post, gateway PLCs, access control units, biometry scanners, webcams, network devices, interphone master stations, intercom servers, interphone client stations, UPS monitoring devices, virtual machine servers, generic windows PCs and generic PLCs. Every device was also reviewed for brand, model, operating system, installed software, and running services, as shown in Table 1.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

INCREASING AVAILABILITY BY IMPLEMENTING SOFTWARE REDUNDANCY IN THE CMS DETECTOR CONTROL SYSTEM

L. Masetti, A. Andronidis, O. Chaze, C. Deldicque, M. Dobson, A. Dupont, D. Gigi, F. Glege, J. Hegeman, M. Janulis, R. Jiménez Estupiñán, F. Meijers, E. Meschi, S. Morovic, C. Nunez-Barranco-Fernandez, L. Orsini, A. Petrucci, A. Racz, P. Roberts, H. Sakulin, C. Schwick, B. Stieger, S. Zaza, CERN, Geneva, Switzerland
P. Zejdl, CERN, Geneva; Fermilab, Batavia, Illinois, USA
U. Behrens, DESY, Hamburg, Germany, O. Holme, ETH Zurich, Switzerland
J. Andre, R. K. Mommsen, V. O'Dell, Fermilab, Batavia, Illinois, USA, G. Darlea, G. Gomez-Ceballos, C. Paus, K. Sumorok, J. Veverka, MIT, Cambridge, Massachusetts, USA
S. Erhan, UCLA, Los Angeles, California, USA, J. Branson, S. Cittolin, A. Holzner, M. Pieri, UCSD, La Jolla, California, USA

Abstract

The Detector Control System (DCS) of the Compact Muon Solenoid (CMS) experiment ran with high availability throughout the first physics data-taking period of the Large Hadron Collider (LHC). This was achieved through the consistent improvement of the control software and the provision of a 24-hour expert on-call service. One remaining potential cause of significant downtime was the failure of the computers hosting the DCS software. To minimize the impact of these failures after the restart of the LHC in 2015, it was decided to implement a redundant software layer for the control system where two computers host each DCS application. By customizing and extending the redundancy concept offered by WinCC Open Architecture (WinCC OA), the CMS DCS can now run in a fully redundant software configuration. The implementation involves one host being active, handling all monitoring and control tasks, with the second host running in a minimally functional, passive configuration. Data from the active host is constantly copied to the passive host to enable a rapid switchover as needed. This paper describes details of the implementation and practical experience of redundancy in the CMS DCS.

ROLE OF THE CMS DETECTOR CONTROL SYSTEM AND EXPERIENCE DURING RUN-I

The DCS of the CMS experiment enables the coherent and safe operation of the detector. It provides the supervision of the experiment, allowing for the operation of the sub-detectors and sub-systems and works in synchronization with the LHC, automatically preparing CMS for data taking whenever LHC is ready to generate collisions. The availability of the DCS is crucial to ensure the operation of the experiment and to maximize data taking time.

The CMS DCS consists of an interconnected federation of independent applications, all of which are based on the WinCC OA control system toolkit from ETM professional. The central CMS DCS team takes responsibility for the infrastructure and platform on which the control applications run. This includes Windows operating system, WinCC OA software and other software dependencies such as the OPC DA servers used to communicate with front-end hardware.

WinCC OA is a modular and distributed software package where the functionality is implemented in independent processes, called *managers*, that communicate via the TCP/IP protocol. The central processing unit is the *event manager* which keeps the current process image in memory and ensures the distribution of data to other managers. The process image is organized in typed, structured process variables, called *data points*. A WinCC OA *project* is a configured set of managers running in a single server, typically connected to a local event manager. A WinCC OA *system* identifies a namespace corresponding to one event manager (or two redundant event managers in case of redundant systems). Many systems can be connected forming a distributed network which enables data exchange between systems.

The CMS DCS ran successfully throughout the first data taking run of the LHC with high availability. However, a potential cause of significant downtime was the failure of server hardware, which would immediately stop the hosted control application. In order to restore full functionality of the CMS control system, the central team would intervene urgently to prepare a new server and target the unavailable application to the new computer. This operation was time consuming (up to 1 hour) and could introduce unexpected problems due to the complexity of the installation process.

In order to mitigate the issue of server failure, it was decided to implement redundancy at the control software layer by running each CMS WinCC OA system on two computers, one located in the underground cavern and the other in the computer centre on the surface, in order to overcome a failure of the services or infrastructure at a single geographical location. In this way, if a single server would fail, all CMS DCS control applications would continue to run and a deferred intervention could then be planned to replace the failed hardware. Details on the selected server hardware can be found in [1]. The present paper presents the final software solutions adopted during the full migration of the CMS DCS to the redundant architecture.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

A DATA ACQUISITION SYSTEM FOR ABNORMAL RF WAVEFORM AT SACLA

M. Ishii#, M. Kago, JASRI/SPring-8, Hyogo 679-5198, Japan
T. Fukui, H. Maesaka, T. Inagaki, T. Maruyama, T. Ohshima, Y. Otake, RIKEN/SPring-8, Hyogo 679-5148, Japan
T. Hasegawa, M. Yoshioka, SPring-8 Service Co., Ltd., Hyogo 679-5165, Japan

Abstract

At the X-ray Free Electron Laser (XFEL) facility, SACLA, we developed a data acquisition (DAQ) system to capture an abnormal radio frequency (RF) waveform for the diagnosis of suddenly occurring failures. The DAQ system consists of VME systems, a cache server, and a NoSQL database system, Apache Cassandra. When the VME system detects an abnormal RF waveform, it collects all the related waveforms of the same shot. The waveforms are stored in Cassandra through the cache server. Before the installation to SACLA, we ensured the performance with a prototype system. We installed the DAQ system into 34 LLRF VME systems at SACLA and six LLRF VME systems at a dedicated accelerator for the SACLA wide-band beam line. The DAQ system successfully acquired rare abnormal waveforms and contributed to the investigation of the failure analysis.

INTRODUCTION

At the X-ray Free Electron Laser (XFEL) facility, SACLA [1], driven with an 8 GeV linac, the stability of the phase and the amplitude of RF in each accelerating structure significantly affect the stability of the electron beam and the FEL intensity. It is quite helpful to collect the shot-by-shot RF data for a feedback control and the diagnosis of a failure source.

We have been using two types of data acquisition systems since the beginning of commissioning in 2011. One is a data logging system with a cycle of several seconds, which provides information for slow fluctuation, such as environmental temperatures and the flow of cooling water [2]. The other is an event-synchronized data acquisition (sync-DAQ) system that collects beam currents, beam positions, and the phase and the amplitude of the RF signals in synchronization with the beam operation cycle at the current maximum of 60 Hz [3]. Shot-by-shot data are tagged with a master trigger number to identify the beam shot to which the data belong. In addition, the sync-DAQ system collects RF waveform data every 10 min. If any fluctuations were observed in shot-by-shot data, it is important for the failure diagnosis to capture the waveform data of when the failure occurred. However, waveform collection every 10 min may not capture an abnormal RF waveform from a rare failure event that may occur only a few times a day.

ishii@spring8.or.jp

In order to collect the abnormal RF waveform data for a failure diagnosis, we developed a data acquisition system to acquire abnormal RF waveforms (abnormal WFM–DAQ system) in 2014. The DAQ system is useful for analyzing the phenomenon, because it has more information than point data, which are sampled from a part of waveform data. In this paper, we describe the installation of the abnormal WFM–DAQ system and an example of the acquired data.

ABNORMAL WFM–DAQ SYSTEM

Low-level RF Control System

The low-level RF (LLRF) control system consists of the 74 RF units of the SACLA linac. Figure 1 shows an overview of an RF unit in the LLRF control system. A trigger delay unit (TDU) board, a DAC board, three ADC boards and a CPU board are installed in a VME crate. The TDU board generates the trigger signals to the ADC and DAC boards, a klystron modulator, etc., whose timing is determined by counting with the 238-MHz master clock. The DAC board generates the pulse modulation pattern signals to the IQ modulator for controlling the RF phase and the RF amplitude. The ADC boards measure the phase and the amplitude of the RF cavity pickup signal from klystron (KLY), pulse compressor cavity, and two accelerating structures, and the high voltage (Vk) and the pulse current (Ik) of the klystron. The TDU board, the DAC board, and the ADC boards synchronize with the 238-MHz master clock.

The ADC board has four input channels. Each channel has four memory banks that can store 512 waveforms data in synchronization with the trigger signal. The data size of

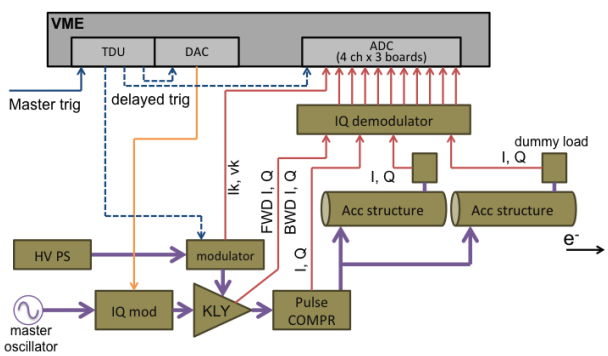


Figure 1: Overview of an RF unit in the LLRF control system.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

DRIVERS AND SOFTWARE FOR MicroTCA.4*

M. Killenberg†, M. Heuer, M. Hierholzer, L. Petrosyan, C. Schmidt, N. Shehzad,
G. Varghese, M. Viti, DESY, Hamburg, Germany
S. Marsching, aquenos GmbH, Baden-Baden, Germany
M. Mehle, T. Sušnik, K. Žagar, Cosylab d.d., Ljubljana, Slovenia
A. Piotrowski, FastLogic Sp. z o.o., Łódź, Poland
T. Kozak, P. Prędki, J. Wychowaniak, Łódź University of Technology, Łódź, Poland

Abstract

The MicroTCA.4 crate standard provides a powerful electronic platform for digital and analogue signal processing. Besides excellent hardware modularity, it is the software reliability and flexibility as well as the easy integration into existing software infrastructures that will drive the widespread adoption of the new standard. The DESY MicroTCA.4 User Tool Kit (MTCA4U) comprises three main components: A Linux device driver, a C++ API for accessing the MicroTCA.4 devices and a control system interface layer. The main focus of the tool kit is flexibility to enable fast development. The universal, expandable PCI Express driver and a register mapping library allow out of the box operation of all MicroTCA.4 devices which are running firmware developed with the DESY board support package. The tool kit has recently been extended with features like command line tools and language bindings to Python and Matlab.

INTRODUCTION

The MicroTCA.4 crate standard [1, 2] provides a platform for digital and analogue data processing in one crate. It is geared towards data acquisition and control applications, providing a backplane with high-speed point to point serial links, common high-speed data buses as well as clock and trigger lines. In typical control applications, large amounts of data have to be digitised and processed in real-time on the front end CPU of the MicroTCA.4 crate.

MTCA4U—The DESY MicroTCA.4 User Tool Kit

The main goal of the DESY MicroTCA.4 User Tool Kit (MTCA4U) [3] is to provide a library which allows efficient, yet easy to use access to the MicroTCA.4 hardware in C++. In addition, it features an adapter layer to facilitate interfacing to control system and middleware software. The design layout of the tool kit is depicted in Fig. 1.

LINUX KERNEL MODULE

The Linux kernel module (driver) provides access to the MicroTCA.4 devices via the PCI Express bus. As the basic access to the PCI Express address space is not device dependent, we follow the concept of a universal driver for all MicroTCA.4 boards. The kernel module uses the Linux Device Driver Model which allows module stacking, so that the driver can be split into two layers: A universal part provides all common structures and implements access to the PCI Express I/O address space. The device specific part implements only firmware-dependent features like Direct Memory Access (DMA), and uses all basic functionality of the universal part. For all devices developed at DESY the firmware will provide a standard register set and the same DMA mechanism, which permits to use a common driver for most boards. For devices from other vendors the universal part enables out-of-the-box access to the basic features, which can be complemented by writing a driver module based on the universal driver part. Like this, the interface in MTCA4U does not change and the new device is easy to integrate into existing software.

The MircoTCA.4 platform allows hot-plugging of all components, which means the PCI Express components can appear and disappear at run time. Usually this is not the case for PCI Express hardware, and not all drivers are prepared to handle this situation. For the drivers provided by MTCA4U special attention has been paid to make the driver hot-plug capable and allow safe operation at all times.

THE C++ DEVICE API

The core piece of MTCA4U is the C++ device access library, which provides a high level interface to the hardware. Its main component is the `Device` class (see Fig. 2), which provides a convenient interface to access registers in the hardware I/O address space. The `Device` class is not accessing the hardware directly but uses the abstract `DeviceBackend` interface, which has several implementations. The `PcieBackend` is accessing the PCI Express hardware through the Linux kernel module, abstracting implementation details like IOCTL sequences from the user. A `DummyBackend` can provide the same set of registers, simulated in the RAM of the CPU module. Deriving from this class, the user can simulate firmware functionality of a specific device, which simplifies unit and integration testing. The latest addition is the `RebotDevice`, which implements the "Register-based over TCP" protocol (ReboT). This allows access to network devices, either via the backplane of the MicroTCA crate or through an external network, broadening the scope of the MicroTCA.4 User Tool Kit also to applications outside the MicroTCA form factor.

* This work is supported by the Helmholtz Validation Fund HVF-0016 "MTCA.4 for Industry".
† martin.killenberg@desy.de

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

SERVICE ASSET AND CONFIGURATION MANAGEMENT IN ALICE DETECTOR CONTROL SYSTEM

M. Lechman[1, 2], A. Augustinus[1], P. M. Bond[1], P. Chochula[1], O. Pinazza[1, 2], A. Kurepin[1, 3]
[1]CERN, Geneva, Switzerland
[2]IP-SAS, Bratislava, Slovakia
[3]INFN-Bologna, Bologna, Italy
[4]RAS/INR, Moscow, Russia

Abstract

ALICE (A Large Ion Collider Experiment) is one of the big LHC (Large Hadron Collider) detectors at CERN. It is composed of 19 sub-detectors constructed by different institutes participating in the project. Each of these subsystems has a dedicated control system based on the commercial SCADA package "WinCC Open Architecture" and numerous other software and hardware components delivered by external vendors. The task of the central controls coordination team is to supervise integration, to provide shared services (e.g. database, gas monitoring, safety systems) and to manage the complex infrastructure (including over 1200 network devices and 270 VME and power supply crates) that is used by over 100 developers around the world. Due to the scale of the control system, it is essential to ensure that reliable and accurate information about all the components – required to deliver these services along with relationship between the assets – is properly stored and controlled. In this paper we will present the techniques and tools that were implemented to achieve this goal, together with experience gained from their use and plans for their improvement.

INTRODUCTION

A Large Ion Collider Experiment (ALICE) [1] is one of the big Large Hadron Collider (LHC) detectors at CERN that is optimized to study the physics of the quark–gluon plasma in nucleus–nucleus collisions. The experiment collaboration consists of 1550 physicists, engineers and technicians from 37 countries.

ALICE is composed of 19 sub-detectors constructed by different institutes participating in the project. Each of these subsystems has a dedicated control software based on the commercial SCADA package "WinCC Open Architecture" [2] and numerous other software and hardware components delivered by external vendors.

The ALICE Controls Coordination (ACC) team supervises the integration of all of these applications into one global Detector Control System (DCS) [3] that allows one to monitor and operate the whole experiment from a central operator console. The group is also responsible for providing shared services (like database, environment monitoring) and for managing complex infrastructure (including over 1200 network devices and 270 VME and power supply crates) that is used by over 100 developers around the world. The DCS includes as well interfaces to external systems maintained by other groups at CERN. Examples of such systems are beam interlock, safety mechanisms or software for acquisition and processing of physics data. The full context of the ALICE DCS is illustrated in Fig. 1.

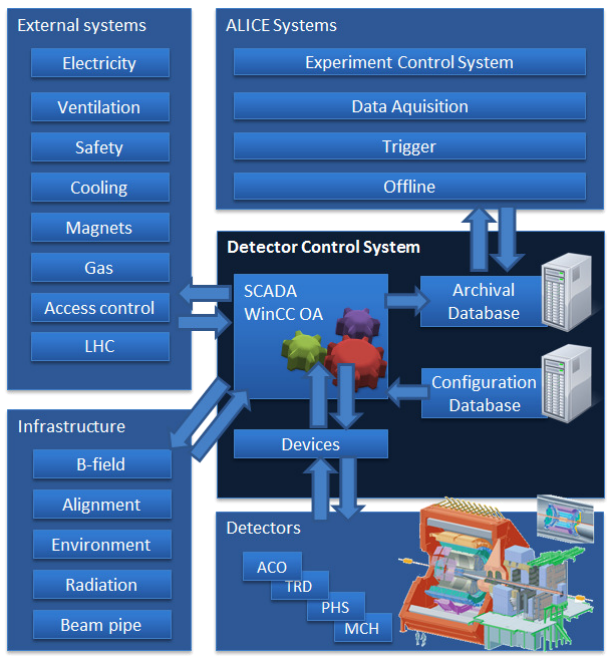


Figure 1: Context of ALICE Detector Control System.

The on-call expert service and training of DCS operators, who supervise the detector from the central console in Run Control Centre, are another tasks of the ACC group.

The team also participates in JCOP (Joint Controls Project) [4] – a collaboration of different experiments at CERN created to share the development effort needed for preparing software components used in their controls systems.

SERVICES AND THEIR ASSETS IN DCS

The aforementioned services require numerous software and hardware assets. Due to the scale of the control system, it is essential to ensure that reliable and accurate information about all these components along with the relationship between them is properly stored and controlled.

To meet these requirements, the central DCS team explores the possibilities to adapt best practises from ITIL (IT Infrastructure Library) [5] to manage its commitments, user groups and resources.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

DATABASE APPLICATIONS DEVELOPMENT OF THE TPS CONTROL SYSTEM

Y. S. Cheng, Jenny Chen, C. Y. Liao, C. H. Huang, P. C. Chiu, Y. T. Chang, K. T. Hsu
National Synchrotron Radiation Research Center, Hsinchu 30076, Taiwan

Abstract

The control system had been established for the new 3 GeV synchrotron light source (Taiwan Photon Source, TPS) which was successful to commission at December 2014. Various control system platforms with the EPICS framework had been implemented and commissioned. The relational database (RDB) has been set up for some of the TPS control system applications used. The EPICS data archive systems are necessary to be built to record various machine parameters and status information into the RDB for long time logging. The specific applications have been developed to analyze the archived data which retrieved from the RDB. One EPICS alarm system is necessary to be set up to monitor sub-system status and record detail information into the RDB if the problem happened. Some Web-based applications with RDB have been gradually created to show the TPS machine status related information. The efforts are described at this paper.

INTRODUCTION

The TPS is a new high brightness synchrotron light source which constructed at the National Synchrotron Radiation Research Center (NSRRC) in Taiwan. It consists of a 150 MeV electron linac, a booster synchrotron, a 3 GeV storage ring, and experimental beam lines. Civil construction had started from February 2010. The construction works had been finished in half of 2013. Accelerator system installation and integration had been started in later 2013. The control system environment was ready in half of 2014 to support final subsystem integration test and commissioning without beam. Commissioning with beam was successful at December 2014.

The EPICS (Experimental Physics and Industrial Control System) [1] is a set of open source software, libraries and applications developed collaboratively and used to create distributed soft real-time control systems for scientific instruments such as the particle accelerators. Many facilities have good practical experiences for the EPICS and adopt it as the accelerator control systems. Many resources and supports are available as well as numerous applications have been developed.

The TPS Control system is also based on the EPICS framework [2]. The EPICS toolkit provides standard tools for display creation, archiving, alarm handling, etc. The big success of EPICS is based on the definition of a standard IOC (Input Output Controller) structure together with an extensive library of driver software for a wide range of I/O cards. The EPICS toolkits which have various functionalities are employed to monitor and to control accelerator system.

The Java-based EPICS data archive system was set up for the TPS project. The Relational Database (RDB) has been used as the data storage mechanism for recording historic EPICS data [3]. The various operation interfaces for browsing archived data have been developed. Taking the performance and redundancy into considerations, the storage servers and database table structures are tuned up relatively.

The Java-based alarm system has been developed as the alarm handler for the TPS control system. A distributed alarm system monitors the alarms in a control system and helps operators to make right decisions and actions in the shortest time. Moreover the historic alarm messages are logged into the RDB system, and the clients can use Java-based graphical operation interface to query and effectively check which system is failure.

Web-based machine status broadcasting is convenient to use with Web browser. The PHP webpage has been created for the TPS machine status broadcasting. The trend data are retrieved from the archive database, and also show variations on the Web-based machine status broadcasting.

The efforts for implementing are summarized as followings.

TPS EPICS DATA ARCHIVE SYSTEM

Software Environment

For the installation and commissioning phases, one EPICS data archive system is necessary to be developed to record various accelerator parameters and machine status information for long time observation. The archive system of CS-Studio [4-5] had been implemented completely to be used as the TPS EPICS data archive system in earlier 2014. The archive engines take PVs data from EPICS IOCs via channel access, and stores them into the data storage.

The CS-Studio based archive system can store data in a relational database (RDB). It applies JDBC libraries [6] in some databases, such as MySQL, Oracle, PostgreSQL, and the archive system includes example database definition (DBD) files to create the required tables for the database dialects. The EnterpriseDB [7] (PostgreSQL) RDB have been used for the EPICS data archive system of TPS project. The PostgreSQL RDB is a good compromise and bigger table sizes. Both the PVs historic data and the Archive Engine configuration are stored in the same relational database. The engine configurations are imported from an XML file format into the database.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

A REDUNDANT EPICS CONTROL SYSTEM BASED ON PROFINET*

Z.Huang, Y. Song, K. Wan, C. Li, G. Liu†, NSRL, USTC, Hefei, Anhui 230029, China

Abstract

This paper will demonstrate a redundant EPICS control system based on PROFIENT. The control system consists of 4 levels: the EPICS IOC, the PROFINET IO controller, the PROFINET media and the PROFINET IO device. Redundancy at each level is independent of redundancy at each other level in order to achieve highest flexibility. The implementation and performance of each level will be described in this paper.

INTRODUCTION

Availability is a key parameter for big scientific facilities, especially for user facilities. High availability drives the reliability demands for the control system. Redundancy is a common approach to improve the reliability and availability of a system. Some efforts are made to implement the system redundancy under EPICS environment, e.g. the redundant IOC is developed based on the redundancy monitor task (RMT) at DESY for the cryogenic system of the European XFEL project [1]; redundancy technology of the IOC is realized via Xen or Linux-HA base on ATCT at Institute of High Energy Physics(IHEP), Chinese Academy of Sciences [2]; the redundant control system is designed at Shanghai Institute of Applied Physics(SINAP), Chinese Academy of Sciences for the Thorium Molten Reactor System(TMSR) project [3].

PROFINET is a standard for Industrial Ethernet, it is defined by PROFIBUS and PROFINET International (PI) and, since 2004, is part of the IEC 61158 and IEC 61784 standards. PI released the Specification "PROFINET IO System Redundancy" on July 19, 2011, which describes the mechanism to build up a redundant PROFINET IO system [4]. There is no product which is consistent with this specification at present. However there are some existing commercial redundancy solutions with PROFINET, e.g. the redundancy system based on Phoenix Redundancy Layer from Phoenix contact [5].

By integrating the commercial solution into EPICS environment, we set up a prototype system. The prototype system consists of 4 levels: the EPICS IOC, the PROFINET IO controller, the PROFINET media and the PROFINET IO device. Redundancy at each level is independent of redundancy at each other level to achieve highest flexibility.

SYSTEM ARCHITECTURE

The system architecture of the redundant control system is shown in Figure 1. The system includes 4 levels: the EPICS IOC, the PROFINET IO controller, the PROFINET media and the PROFINET IO device.

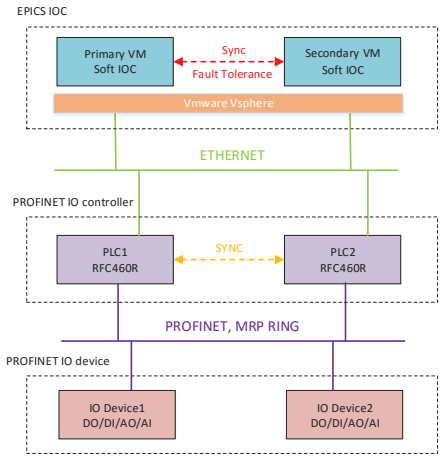


Figure 1: System architecture of the redundant control system.

EPICS IOC

VMware vSphere is the industry's leading and most reliable virtualization platform. It achieves always-available IT with live migration for virtual machines and high availability for applications in virtual machine (VM) clusters. VMware provides 4 features to improve availability: VMware vMotion, VMware Storage vMotion, VMware High Availability (HA), VMware Fault Tolerance (FT) [6]. In this redundant system, VMware FT is chosen to achieve hot standby for the EPICS IOC.

VMware FT provides continuous availability for VMs by creating and maintaining a Secondary VM that is identical to, and continuously available to replace, the Primary VM in the event of a failover situation. The Primary and Secondary VMs continuously exchange heartbeats. This allows the virtual machine pair to monitor the status of one another to ensure that Fault Tolerance is continually maintained [7].

The EPICS softIOC running on the VM communicates with the redundant pair of two RFC 460R PLCs via TCP/IP. An EPICS driver has been developed for the communication between the EPICS softIOC and the RFC 460R PLCs [8].

PROFINET IO Controller

The RFC 460R PLC from Phoenix Contact is a high performance controller that has been extended to offer redundancy functionality. The redundant pair consists of two synchronized RFC 460R PLCs connected via fiber optics. The built-in fiber optic interface is used for synchronization and adjustment between the connected devices [9].

* Work supported by National Natural Science Foundation of China (11375186)

† gfliu@ustc.edu.cn

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

DESIGN OF CONTROL NETWORKS FOR CHINA INITIATIVE ACCELERATOR DRIVEN SYSTEM

Zhiyong He*, Yuxi Luo, Yuhui Guo, Qiang Zhao, Wenjuan Cui, Yichuan He
Institute of Modern Physics, Chinese Academy of Sciences, Lanzhou, China

Abstract

In this paper, we report the conceptual design of control networks used in the control system for China initiative accelerator driven sub-critical (ADS) facility which consists of two accelerator injectors, a main accelerator, a spallation target and a reactor. Because different applications have varied expectations on reliability, latency, jitter and bandwidth, the following networks have been designed for the control systems, i.e. a central operation network for the operation of accelerators, target, and reactor; a reactor protection network for preventing the release of radioactivity to the environment; a personnel protection network for protecting personnel against unnecessary exposure to hazards; a machine protection network for protecting the machines in the ADS system; a time communication network for providing timing and synchronization for three accelerators; and a data archiving network for recording important measurement results from accelerators, target and reactor. Finally, we discuss the application of high-performance Ethernet technologies, such as Ethernet ring protection protocol, in these control networks for CIADS.

INTRODUCTION

China is the world's most populous country with a fast-growing economy that has led it to become the largest energy consumer in the world in 2010. In China, nuclear power accounts for relatively small shares (nearly 1%) of the country's total energy consumption in 2011. China is actively promoting nuclear power as a clean, efficient, and reliable source of electricity generation. Driven by the national demand for safe disposal of nuclear waste as well as the potentials for advanced power generation, the Chinese Academy of Sciences initiated an accelerator driven sub-critical (ADS) program in 2011 under the frame of "Strategic Priority Research Program"[1]. The ultimate goal of the China ADS program is to build an industrial demo facility for ADS technology.

In an ADS system, a heavy metal spallation target located at the centre of a sub-critical core is bombarded by the high-energy protons from an accelerator. Figure 1 shows a drawing of China ADS facility which consists of an accelerator, a spallation target and a sub-critical reactor. As shown in Figure 1, a LINAC accelerator which includes two injectors, a medium energy beam transport line (MEBT), several spokes, and a high energy beam transport line (HEBT) is used in China ADS facility. In this paper, we report the conceptual design of control networks used in China ADS facility.

* corresponding author, Email: zyhe@impcas.ac.cn
This work is supported by Strategic Priority Research Program of Chinese Academy of Sciences (XDA03010000).

NETWORKS IN CHINA ADS FACILITY

Because different applications have varied expectations on reliability, latency, jitter and bandwidth, several networks have been used in ITER (International Thermonuclear Experimental Reactor) (e.g. [2]-[3]). The China ADS system includes two totally different facilities, accelerator and reactor. To successfully integrate the two different facilities into one system, several networks are required in the control system for China ADS system, while some networks are designed especially for accelerator and other networks are used especially for reactor.

Central Operation System and Network

Control systems in ADS may include a central control system and several local control systems. The central control system is used to control the overall ADS facility, in particular, to exactly couple the high-energy beam from the accelerator to the spallation target located at the centre of the reactor core. The central operation network (CON) used in the central control system provides interface between the operation system in the central control room and various controller and sensors in the system.

Reactor Protection System and Network

To guarantee the integrity of the reactor and to avoid an undue risk to the public health and safety, the plant design incorporates a reactor protection system (RPS). The overall purpose of the RPS is to prevent the release of radioactivity to the environment. To meet this objective, this system is capable of supplying reactor and component trip signals if safe operating limits are exceeded, and initiates the engineered safety features actuation(s) if an accident occurs.

The RPS contains two complete and independent trains of monitoring systems. If a monitoring system detects an unsafe event, signals are transmitted via the reactor protection network (RPN) to the RPS. The RPS determines whether the coincidence for a reactor trip function is satisfied. If so, the protection system opens the reactor trip breakers. Opening these breakers removes power from the control rod drive mechanisms, allowing the rods to fall into the reactor core. If an accident occurs and an engineered safety features actuation is required, the protection system actuates the appropriate safety equipment.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

CONTROLLING CAMERA AND PDU

O.J. Mokone, T. Gatsi, SKA South Africa NRF, South Africa

Abstract

The 64-dish MeerKAT radio telescope, currently under construction in South Africa, will become the largest and most sensitive radio telescope in the Southern Hemisphere until integrated with the Square Kilometre Array South Africa (SKA SA). This paper will present the software solutions that the MeerKAT Control and Monitoring (CAM) team implemented to achieve control (pan, tilt, zoom and focus) of the on-site video camera using the Pelco D protocol. Furthermore this paper will present how the outlets of the Power Distribution Unit (PDU) are switched on and off using Simple Network Management Protocol (SNMP) to facilitate emergency shutdown of equipment.

INTRODUCTION

The current running KAT-7 radio telescope (a MeerKAT prototype), which is running in the Karoo, is utilising three fixed cameras. Fixed cameras sometimes do offer appreciable performance in terms of resolution and image quality but do have amongst others a host of shortcomings for all their intents and purposes.

Despite them being more durable and needing less of mechanical service, the fixed cameras have the disadvantage of not covering a large area in the vicinity of the dishes and hence increases the cost of purchasing more fixed cameras in order to monitor the telescopes and the surrounding area.

To avert this challenge, the SKA SA MeerKAT team procured a Pan, Tilt and Zoom (PTZ) video camera. This Camera has been incorporated into the MeerKAT system because it can closely monitor the area under surveillance through the exploitation of the zoom-pan-tilt functionalities that are built within the camera.

The PTZ camera implements a state of the art technology, which is more appropriate for the MeerKAT project as it surveils a wide area around the dishes and allows for the pan, tilt and zoom capabilities. The PTZ camera is connected to the Power Distribution Unit (PDU). The PDU makes it possible to switch on/off the camera electronics that are connected to the specific PDU outlets remotely. The software that control and monitor camera was developed using Python programming language running on Ubuntu 14.04 LTS.

CAMERA SOFTWARE IMPLEMENTATION

Karoo Array Telescope Control Protocol (KATCP) [1] is a simple ASCII communication protocol layered on top of TCP/IP and is used for the monitoring and control of hardware devices. One of the most prevailing features of the KATCP protocol is its capability to interrogate KATCP servers for sensors. Interrogation of sensors [2] and requests, down to device level through KATCP, supports fluid run-time detection of system configuration, e.g. when a new sensor is added to any level of the system (including a hardware device), the rest of the CAM automatically discovers the change on-line on, and includes it in the monitor store and in its interfaces. Based on the functionality of the KATCP protocol, the video camera software had to implement similar technologies that respond to KATCP command messages.

Pelco D is a popular PTZ camera control protocol used in the CCTV industry. The camera control software uses Pelco D protocol over RS485, a TCP/IP physical interface. This is a client-server architecture, where the server is KATCP device server that connects directly to the camera device and the client is the CAM (Control And Monitoring) system software. One of the major roles of the server is to serve as a protocol translator; it translates Pelco D command messages to KATCP command messages and vice versa. Pelco D consists of 7 hexadecimal bytes. Table 1 (below) shows Pelco D message formats and their descriptions.

Table 1: Pelco D Message Format

Byte Number	Description
Byte 1	Synch Byte
Byte 2	Address
Byte 3	Command 1
Byte 4	Command 2
Byte 5	Data 1
Byte 6	Data 2
Byte 7	Checksum

Byte 1 is the synchronization byte, which is always fixed to FF. Byte 2 is logical address of the camera being controlled. Byte 7 is checksum, which is the 8 bit (modulo 256) sum of the payload bytes byte 2 through byte 6 in the message. The rest of the byte numbers depends on the commands that one is passing to the server.

PRESETS

The PTZ Camera is capable of storing and accessing up to 127 presets. These stored presets reside in a non-volatile memory within the camera. They can be accessed using an identity number from 1 through 127. The identity that is used to store a preset can subsequently be used to “go-to” to this preset, thus causing the camera to move to the previously stored configuration of pan, tilt,

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

INTERFACING EPICS TO THE WIDESPREAD PLATFORM MANAGEMENT INTERFACE IPMI

M. Ritzert*, Heidelberg University, Germany

Abstract

The Intelligent Platform Management Interface (IPMI) is a standardized interface to management functionalities of computer systems. The data provided typically includes the readings of monitoring sensors, such as fan speeds, temperatures, power consumption, etc. It is provided not only by servers, but also by µTCA crates that are often used to host an experiment's control and readout system. Therefore, it is well suited to monitor the health of the hardware deployed in HEP experiments. In addition, the crates can be controlled via IPMI with functions such as triggering a reset, or configuring IP parameters. We present the design and functionality of an EPICS module to interface to IPMI that is based on ipmitool. It supports automatic scanning for IPMI sensors and filling the PV metadata (units, meaning of status words in mbbi records) from the IPMI sensor information. Most importantly, the IPMI-provided alarm thresholds are automatically placed in the PV for easy implementation of an alarm system to monitor IPMI hardware.

IPMI

The Intelligent Platform Management Interface (IPMI) is a common interface to monitoring and management data of servers, ATCA and µTCA crates and similar devices. Monitoring data is provided as analog or digital sensor readings. The sensor metadata provided includes several alarm thresholds that can be used in a monitoring system. Server management is performed via messages sent to the IPMI device from the control PC.

The structure of a typical IPMI system in an ATCA crate is shown in Fig. 1 The typical mode of access to IPMI systems in a distributed system landscape is via Ethernet LAN. The connection is received on the server side by the IPMI baseboard management controller (BMC) that operates independently of the host's CPU, often even when the host device is currently powered down. The BMC communicates via the intelligent platform management bus (IPMB) with other management controllers (MCs) using a modified I²C protocol. Each MC manages a set of sensors, and can be queried to enumerate and read these sensors in a standardized format. Automatic detection of the sensors available in a system is therefore possible, and all sensors are accessed in the same way; all details of the actual hardware access are hidden by the MC.

Since IPMI access is available also when the host device is powered off, it is also a useful tool to remotely manage compute farms. Servers or crates can be reset, started, or stopped without invoking the OS.

When custom hardware is built for IPMI-enabled platforms such as ATCA crates, the IPMI system can be extended by adding new MCs. The required connections to interface to the IPMB are provided on the backplane's connectors.

IPMITOOL-IOC

Ipmitool [1] has been chosen as the basis to develop an interface (input/output controller, IOC) from EPICS to IPMI systems. The main functionality of ipmitool is linked into a static library that is not installed into the target system. Therefore, the source code of ipmitool must be available at build time.

Since ipmitool is used for the communication, advanced features of IPMI such as authentication are easily implemented. Compatibility with a wide range of hardware is given; the systems examined during development include ATCA crates and 19" servers of several brands.

The IOC is written in C++, interfacing to the EPICS CA libraries and the ipmitool library written in plain C. Since data queries via IPMI are "slow" in the EPICS sense, the asynchronous PV support of EPICS is used. Typical scanning times for all sensors in a system could be in the order of a few seconds.

A custom logging library written in C++ is used to send textual log messages from the IOC to a STOMP server that can pass them on to JMS2RDB, so that they can eventually be displayed in the "Message History" view of CSS. This way, all deployed IOCs report their messages to a central place, and messages are less likely to be overlooked.

Usage of the Module

First, the connection to the IPMI host has to be established. From the IOC, this can be done by calling the new ipmiConnect function. It accepts parameters to define a numeric handle for the connection, the host name, and if required, the user name, password, and access level. Next, the device is scanned for sensors by calling ipmiScanDevice. The result of the scan is used to populate the meta data of the PVs that are loaded via the standard dbLoadRecords call.

Optionally, the result of the device scan can be dumped in EPICS database format to serve as a starting point when preparing the IOC to control the device.

The link addresses in the database use the AB_IO format "#Lw Ax Cy Sz @p", where w is the id of the link established with ipmiConnect, x is the IPMB target address, and y is the sensor id. z is reserved for future use (possibly for IPMB bridging addresses), as is p.

Automatic Sensor Detection

The IOC uses ipmitool's functions to enumerate the sensors in a system. It maintains a work list of IPMB addresses

* michael.ritzert@ziti.uni-heidelberg.de

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

DATA DRIVEN SIMULATION FRAMEWORK

Amar Banerjee, Subhrojyoti Roy Chaudhuri, Puneet Patwari, TRDDC, Pune, India
Lize Van Den Heever, SKA South Africa, Cape Town, South Africa

Abstract

Control systems for radio astronomy projects such as MeerKAT[1] require testing functionality of different parts of the Telescope even when the system is not fully developed. Usage of software simulators in such scenarios is customary. Projects build simulators for subsystems such as dishes, beam-formers and so on to ensure the correctness of a)their interface to the control system b)logic written to coordinate and configure them. However, such simulators are developed as one-offs, even when they implement similar functionality. This leads to duplicated effort impacting large projects such as Square Kilometer Array[2]. To mitigate this we leverage the idea of data driven software development and conceptualize a simulation framework that reduces the simulator development effort to: 1)capturing all the necessary information through instantiation of a well-defined simulation specification model 2)configuring a reusable engine that performs the required simulation functions based on the instantiated and populated model provided to it as input. We discuss the results of a PoC for such a simulation framework implemented in the context of Giant Meter-wave Radio Telescope[3] in this paper.

INTRODUCTION

Large projects that involve implementation of large hierarchy of control systems generate dependencies across the controllers to be developed and as a result on the teams developing them. This inter-dependency creates problems in the verification of controllers developed across teams since the individual teams follow their own time-lines which often are not well synchronized.

A general solutions to this problem is replacing the missing components with simulators to aid the verification of the dependent components. This however increases the effort to manually develop the simulators which also at times results in duplication of efforts.

Although the need for simulators in such scenarios seem to be essential, it is desired to reduce the cost of building such simulators since it might impact the overall cost of projects such as Square Kilometer Array(SKA) significantly, as such projects have large number of modules.

With the proposed simulation and testing framework it becomes possible to achieve the goal of verification of the module with simulation, following the model driven approach. The framework incorporates meta models of the controller node, the simulator and testing based on which it automatically generates the simulators. Much of this information is derived from the Self Description Data (SDD) which contains description of the Controller to be developed.

In this paper we discuss the simulator and test framework, the architecture and the working of the framework.

The paper starts with a discussion on standard practice involved in control system testing and verification using MeerKAT as the Case study. It is followed by the architecture section where we describe our understanding and design of the framework. In the final section we present a proof of concept showing the use of our simulation framework. This is followed by a conclusions section where we describe the conclusions of our work.

STANDARD PRACTICE

MeerKAT Case Study

One of the development practices of the MeerKAT CAM (Control and Monitoring) team was to use a fully simulated system at all times. The MeerKAT CAM team has been using simulators extensively and continuously for development, testing and qualification of the CAM subsystem functionality throughout the MeerKAT project life-cycle, since the early days of Fringe Finder for the very first two antennas in the Karoo, through KAT-7[4] to this day for MeerKAT with each array release. It is possible to run a MeerKAT CAM configuration including only simulated devices, or any combination of real and simulated devices combined. This allows full software development, unit testing, integration testing, and CAM subsystem qualification without any dependency on the hardware being available.

While the CAM team was responsible for developing most of the simulators, some of these device simulators were contractually delivered by the subsystem contractor to ensure that, given their knowledge of the device, the behaviour of the device is reflected with sufficient accuracy by the device simulator. In cases where the subsystem contractor did not deliver such a simulator, the CAM team developed a software simulator for the KATCP interface. Preparing the simulators gave the CAM team a valuable opportunity to gather information about the behaviour of the other subsystems even before those subsystems have been fully developed and are ready for integration.

Each simulator represents the specific messages on KATCP (KAT Control Protocol) interface for a subsystem (commands/requests and monitoring points/sensors), as

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

A SELF-CONFIGURABLE SERVER FOR CONTROLLING DEVICES OVER THE SIMPLE NETWORK MANAGEMENT PROTOCOL

V. Petrosyan, V. Rybnikov, DESY, Hamburg, Germany

Abstract

The Simple Network Management Protocol (SNMP) is an open-source protocol that allows manufacturers to utilize it for controlling and monitoring their hardware. More and more SNMP-manageable devices show up on the market that can be used by control systems for accelerators. SNMP devices are also being used at the free-electron laser (FLASH [1]) and planned to be used at the European X-ray Free Electron Laser (XFEL[2]) at DESY, Hamburg, Germany. To provide an easy and uniform way of controlling SNMP devices a server program has been developed. The server configuration, with respect to device parameters to control, is done during its start-up and driven by the manufacturer Management Information Base (MIB) files provided with SNMP-enabled hardware. The list of controlled parameters can also be tailored by user's configuration files.

The SNMP server integrates the various type devices into the Distributed Object-Oriented Control System (DOOCS [3]) used at FLASH and XFEL. Several SNMP devices of different types can be controlled by one SNMP server simultaneously.

SNMP AT A GLANCE

SNMP is an application-layer protocol defined by the Internet Architecture Board (IAB) in RFC1157 (May of 1990) [4]) for exchanging management information between network devices. It is a part of Transmission Control Protocol/Internet Protocol (TCP/IP [5]) suite. SNMP uses User Datagram Protocol (UDP) [6] for the transport layer. SNMP requests are transmitted as UDP datagrams over a connectionless transmission interface between SNMP manager and SNMP agent. Management Information Base (MIB) is a database containing information about elements to be managed. It represents Manageable Objects (MO) as resources. Every MO (device attribute) has a unique Object Identifier (OID) that can be used by queries. It also provides a map between numeric OIDs and a human-readable text.

SNMP agent is a daemon process running on a network device. It exposes device attributes to managers. SNMP managers control the attributes via Get/Set requests. Managers can have access to those attributes that are allowed by the agent. SNMP agents can also provide traps (asynchronous messages). Traps are mainly used for significant events in the device (e.g. power-off).

SNMP SERVER IMPLEMENTATION

The SNMP server is implemented as a DOOCS server.

It allows to integrate SNMP devices into FLASH and XFEL control systems in a transparent way. The DOOCS SNMP server as an SNMP manager provides access to attributes of the controlled device by means of DOOCS properties.

Server Design

The server uses the NET-SNMP library [7] that provides all functionality required to work with SNMP agents. The library allows:

- to retrieve the list of attributes controlled by an SNMP agent,
- read and write attributes values.

The list of attributes can be acquired starting from a certain node of the MIB tree. This node is defined by one DOOCS property of the SNMP server. Once the full list of attributes is available for the SNMP server it filters out the attributes of interest and creates for every selected SNMP attribute a set of DOOCS main (see Table 1) and optional (see Table 2) properties. Every property name in the set starts with the SNMP attribute name with upper case letters. The filtering of attributes is implemented via white or black lists provided by users.

Reading/writing an SNMP attribute value is implemented via a dedicated DOOCS property. There are two ways to read the attribute value:

- periodically via the server internal update timer,
- immediately on access to the DOOCS property.

Writing the attribute value is done on the corresponding DOOCS property (NAME) change.

Table 1: Main DOOCS Properties for SNMP Attribute

Name	Purpose
NAME.MIB	MIB attribute name
NAME.DESC	Attribute XML description
NAME.ACCESS	Access mode (e.g. read-only)
NAME.TYPE	Data type (e.g. ASN_INTEGER)
NAME	Value of the attribute
NAME.SL	List of valid values to write

According to “Structure of Management Information Version2” (SMIv2, RFC2578 [8]) the data types used by SNMP agents don't include floating point or double type.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

HIGH LEVEL SOFTWARE STRUCTURE FOR THE EUROPEAN XFEL LLRF SYSTEM

Christian Schmidt*, Valeri Ayvazyan, Julien Branlard, Łukasz Butkowski, Olaf Hensler, Martin Killenberg, Mathieu Omet, Sven Pfeiffer, Konrad Przygoda, Holger Schlarb, DESY, Hamburg, Germany
Adam Piotrowski, FastLogic sp. z o.o., Łodz, Poland
Wojciech Cichalewski, Filip Makowski TUL-DMCS, Lodz, Poland

Abstract

The Low Level RF system for the European XFEL is controlling the accelerating RF fields in order to meet the specifications of the electron bunch parameters. A hardware platform based on the MicroTCA.4 standard has been chosen to realize a reliable, remotely maintainable and high performing integrated system. Fast data transfer and processing is done by field programmable gate arrays (FPGA) within the crate, controlled by a CPU via PCIe communication. In addition to the MicroTCA.4 system, the LLRF comprises external supporting modules also requiring control and monitoring software. In this paper the LLRF system high level software used in E-XFEL is presented. It is implemented as a semi-distributed architecture of front end server instances in combination with direct FPGA communication using fast optical links. Miscellaneous server tasks have to be executed, e.g. fast data acquisition and distribution, adaptation algorithms and updating controller parameters. Furthermore the inter-server data communication and integration within the control system environment as well as the interface to other subsystems is described.

INTRODUCTION

The Deutsches Elektronen-Synchrotron (DESY) in Hamburg is currently building the European X-ray Free Electron Laser (E-XFEL) [1]. This hard X-ray light source generates up to 27000 coherent laser pulses per second with a duration of less then 100 fs and a wavelength down to 0.05 nm. For this, electrons have to be accelerated to 17.5 GeV using a 2 km particle accelerator based on superconducting radio frequency technology. Precision regulation of the RF fields inside the accelerating cavities is essential to provide a highly reproducible and stable electron beam. RF field regulation is done by measuring the stored electromagnetic field inside the cavities. This information is further processed by the feedback controller to modulate the driving RF source, using a low level RF system. Detection and real-time processing are performed using most recent FPGA techniques. Performance increase demands a powerful and fast digital system, which was found with the Micro Telecommunications Computing Architecture (MicroTCA.4) [2]. Depending on the application either parallel processing for latency efficiency or pipelined processing of a large number of similar RF signals is aspired. In-tunnel installation with limited access requires a compact system with implied redundancy. Further remote access to status information, software upgrades and maintenance is ensured. Finally the modularity and scalability of this system guarantees to have later upgrades in order to meet demands within a lifetime of the machine and beyond. DESY currently is operating the Free Electron Laser (FLASH), which is a user facility of the same type as E-XFEL but at a significantly lower maximum electron energy of 1.2 GeV. The LLRF system for FLASH is equal to the one of E-XFEL, which allows for testing, developing and performance benchmarking in advance of the E-XFEL commissioning [3]. A picture of the FLASH installation can be found in Fig. 1.



Figure 1: Installation of MicroTCA crates inside the FLASH tunnel underneath the accelerator module. The radiation shielding (yellow cabinet) protects the electronic rack inside. The insert shows the MicroTCA crate installed in this rack.

LLRF SYSTEM LAYOUT

The present LLRF system for a single E-XFEL RF station controlling 32 cavities is given in Fig. 2. LLRF systems with a different quantity of cavities are build as a subset of this configuration. Machine operation is performed in a pulsed mode. The duty cycle, i.e. the repetition rate (10 Hz) of the RF pulse to the RF pulse length (about 1 ms) is about 1 %. The impact to the RF regulation is such that within in a regular mode of unchanged operating conditions, i.e. stable working point consecutive pulses are similar, however the initial conditions vary as well as long term changes occur. In

* christian.schmidt@desy.de

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

THE EPICS ARCHIVER APPLIANCE

Murali Shankar, Luofeng Li, SLAC, Menlo Park, CA, USA*
Michael Davidsaver, BNL, Upton, New York, USA
Martin Konrad, National Superconducting Cyclotron Laboratory, Michigan State University, East Lansing, MI 48824, USA**

Abstract

The EPICS Archiver Appliance was developed by a collaboration of SLAC, BNL and MSU to allow for the archival of millions of PVs, mainly focusing on data retrieval performance. It offers the ability to cluster appliances and to scale by adding appliances to the cluster. Multiple stages and an inbuilt process to move data between stages facilitate the usage of faster storage and the ability to decimate data as it is moved. An HTML management interface and scriptable business logic significantly simplify administration. Well-defined customization hooks allow facilities to tailor the product to suit their requirements. Mechanisms to facilitate installation and migration have been developed. The system has been in production at SLAC for about 2 years now, at MSU for about 2 years and is heading towards a production deployment at BNL. At SLAC, the system has significantly reduced maintenance costs while enabling new functionality that was not possible before. This paper presents an overview of the system and shares some of our experience with deploying and managing it at our facilities.

REQUIREMENTS

At SLAC, BNL and MSU, we use EPICS as our control system for our various facilities; many of these have several million Process Variables (PVs) that are used for monitoring and control. A common requirement is to be able to archive some of these PV's to facilitate troubleshooting and analysis.

Archiving a million PVs has some non-technical requirements. In addition to scaling gradually by adding additional appliances/hardware, we need better support for managing all aspects of the system. This includes establishing policies for archiving and the flexible configuration of archiving on a per PV basis. The ability to manage the system using a web based UI and from within scripts is a necessity. These management functions must include the ability to make changes to the PV archiving configuration without having to restart the system. Finally, we needed a simple migration path from the ChannelArchiver [1, 2].

In terms of storage, the data rates are high enough to require faster storage but the data volumes are also large. In addition, the archivers require the storage to be always available. A solution using multiple stages of storage is most economical. The first stage is expensive fast storage local to the appliance and thus always available. The final stages of storage can be much slower and cheaper with relaxed requirements for availability including support for tape storage. The final design should be extensible enough to support storage from external vendors like S3.

Our customers also wanted a solution that focused on data retrieval performance. The most common use of archive data is to help trouble shoot machine operations. Thus, the bulk of the data retrieval requests are for recent data. Data can then be moved off to cheaper, slower storage once it has aged. Our goal was to be able to retrieve a day's worth of 1Hz double data in less than 0.5 seconds; the actual retrieval numbers are much faster. In addition, support for various forms of data reduction during data retrieval is critical.

ARCHITECTURE

The EPICS Archiver Appliance uses an appliance model for deployment. An installation is a cluster of appliances. Each appliance (see Fig. 1) has multiple storage stages and multiple processes.

A wide variety of storage configurations are possible (on a per PV basis). The out-of-the-box configuration has these storage stages

- STS (Short term store) - The most recent couple of hours worth of data is typically stored here. This is typically a RAM disk.
- MTS (Medium term store) - The most recent couple of days worth of data is stored here. At SLAC, we use RAIDed 15k SAS drives for the MTS. At MSU, these are mirrored 1.2TB SAS drives.
- LTS (Long term store) - The rest of the data is stored here. At SLAC, this is bulk storage (with tape backups) that we rent from our computing dept. This is a GPFS filesystem located elsewhere and is mounted over NFS. At MSU, this is a NetApp alliance with 2.8 TB of storage.

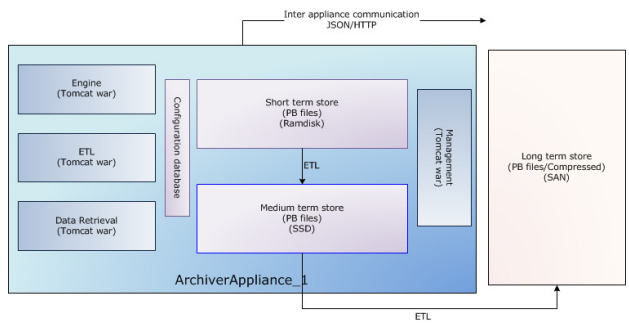


Figure 1: Components of an appliance.

* Work is supported by the U.S. Department of Energy, Office of Science under Contract DE-AC02-76SF00515 for LCLS I and LCLS II
** This work was supported in part by the National Science Foundation under the Cooperative Agreement PHY-11-06007

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

THE EVOLUTION OF THE SIMULATION ENVIRONMENT IN ALMA

T. Shen, R. Soto, N. Saez, G. Velez, S. Fuica, T. Staig, N. Ovando, Joint ALMA Observatory, Santiago, Chile
J. Ibsen, European Southern Observatory, Santiago, Chile

Abstract

The Atacama Large Millimeter /sub millimeter Array (ALMA) has entered into operations transition phase since 2014. This transition changed the priorities within the observatory. Most of the available time will be dedicated to science observations instead of technical time for commissioning activities including software testing. The lack of the technical time surfaces one of the weakest points in the existent infrastructure available for software testing: the simulation environment of the ALMA software. The existent simulation focuses on the functionality aspects but not on the real operation scenarios with all the antennas. Therefore, scalability and performance problems introduced by new features or hidden in the current accepted software cannot be verified until the actual problem explodes during operations. Therefore, it was planned to design and implement a new simulation environment, which must be comparable, or at least, be representative of the production one. In this paper we will review experiences gained and lessons learnt during the design and implementation of the new simulated environment.

OVERVIEW

The Atacama Large Millimeter /sub millimeter Array (ALMA) started the operations phase transition in 2014. This transition changed the priorities within the observatory. Most of the available time will be dedicated to science observations instead of technical time for commissioning activities including software testing. (since software/hardware integration was the former priority during construction phase). The lack of the technical time surfaces one of the weakest points in the existent infrastructure available for software testing: the simulation environment of the ALMA software. In ALMA, simulation capabilities were initially developed to satisfy Control and Correlator subsystem needs, supplying them with virtual hardware devices to interface with the software components being developed. Additional simulation layers and capabilities were added during the years by different teams, but they were focused on the functionality aspect but not on the real operation scenarios with the whole array. Therefore, scalability and performance problems introduced by new features or hidden in the current accepted software cannot be verified until the actual problem explodes during operation. The lack of a representative testing environment will seriously impact the efficiency of the ALMA incremental software release process.

It was planned to design and implement a new simulation environment, which must be comparable, or at least, be representative of the production one. Duplicating the production environment was not an option, since it would be prohibitive from the point of view of the associated cost. Adjustments in the existent simulation architecture had to be introduced, but with special care on keeping the simulation environment comparable in terms of CPU load, network bandwidth throughput, memory usage and software configurations. The selected platform to provide computing power is based on blade technology of Cisco Unified Computing System (UCS). The new simulation platform will provide the required amount of time for testing purposes, at same time it allows us to maximize the efficiency of the reduced technical time available with operational hardware, which will be dedicated only for the final validation of a new release and small set of features that interact directly with the system. In this paper we will review experiences gained and lessons learnt during the design and implementation of a new simulation environment.

PRODUCTION ENVIRONMENT

The hardware related to the ALMA array (antennas, photonic references, correlators, etc) are controlled through an STE [1], which is a collection of servers, real time computers, network switches, storages and databases, that are configured accordingly as a single platform to support the execution of the ALMA observing software [2]. In the Fig. 1, the list of servers, and real time computers are shown.

EXISTENT SIMULATION ENVIRONMENT

The existing simulation platform is a reduced scale version of the STE dedicated for operation, which contains few servers are available to simulate an array of few antennas. This simulation environment has been very useful during the AIV stage [3] where the focus was the commissioning of newly assembled antennas. At that moment, no more than two antennas were usually configured into a single STE.

It became evident that the existent simulation environment needed to be upgraded since the pressure to use the production hardware for scientific observation is extremely high. Currently, the most critical deficiency in the current simulation environment is in the deployment aspects. It differs in terms of network layout, network throughput, computing power, available memory and disk I/O.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

EPICS PV MANAGEMENT AND METHOD FOR RIBF CONTROL SYSTEM

A. Uchiyama#, M. Komiyama, N. Fukunishi
RIKEN Nishina Center, Wako, Saitama, Japan

Abstract

For the RIKEN Radioactive Isotope Beam Factory (RIBF) project, the Experimental Physics and Industrial Control System (EPICS)-based distributed control system is utilized in Linux and vxWorks. Utilizing network attached storage (which has a high-availability system) as a shared storage, common EPICS programs (Base, Db, and so on) are shared by each EPICS Input/Output Controller (IOC). From the initial development of RIBF control system, it has continued to grow and consisted of approximately 50 EPICS IOCs and more than 100,000 EPICS records. Because RIBF has been constructed by extending RIKEN Accelerator Research Facility (RARF) in a previous project, the controllers for RARF are also utilized for RIBF control system. In this case, the dependence between the EPICS records and EPICS IOCs becomes complicated. For example, it is not easy to know the accurate EPICS record name information using only the device information. Therefore, we constructed a new management system for the RIBF control system to easily call up the detailed information. In the system, by parsing startup script files (st.cmd) to run EPICS IOCs, all EPICS records and EPICS fields are stored in the PostgreSQL-based database. By utilizing these stored data, we succeeded in developing Web-based management and search tools.

INTRODUCTION

The RIKEN Radioactive Isotope Beam Factory (RIBF) accelerator facility consists of five cyclotrons, including a superconducting ring cyclotron, and two linear accelerators as injectors [1]. For the RIBF, we constructed a control system based on the Experimental Physics and Industrial Control System (EPICS) for magnet power supplies, beam diagnostic instruments, vacuum control systems, and so on [2]. Different types of devices are connected via the EPICS Input/Output Controllers (IOCs) as front-end controllers for accelerator operation. For example, VME-based IOCs are adapted for magnet power supplies, and beam diagnostic and vacuum systems (e.g., Faraday cup and beam profile monitor) are constructed using Linux-based IOCs connected with N-DIMs, which are Ethernet-based control devices originally developed by the RIKEN Nishina Center [3]. Additionally, CAMAC, GPIB, Programmable Logic Controller (PLC), and others are also utilized as control devices with EPICS for the various applications in RIBF control system. The types of control devices and EPICS IOCs are listed in Table 1.

Currently, the EPICS-based RIBF control system consists of 51 IOCs and approximately 110,000 EPICS records. We have constructed the EPICS runtime database using an Oracle-based database. In this case, using the Oracle database has some advantages from the viewpoint of management because it is available to automatically make EPICS runtime database files easily using a program. In case of a large number of EPICS runtime database files for the same purpose, the abovementioned approach is efficient. Since RIBF control system has been constructed by extending the control system of RIKEN Accelerator Research Facility (RARF), which is the previous project, the relationship among their controllers are complicated. Therefore, we often manually develop EPICS runtime database files without using Oracle-based database.

For system maintenance and development in this situation, we need to identify the IOC connected to the process variables (PVs) because the RIBF control system does not use the UDP broadcasts for *ca_search* from EPICS clients on the subnet. However, the relationship between the IOCs and the PVs is confused, and an efficient system environment was not provided for developers and operators. Therefore, we have constructed a system to manage the EPICS PV to solve the complication.

Table 1: Current Status of the Type of EPICS IOC in the RIBF Control System

IOC Platform	Connected Control Device	Type	Number of IOCs
Linux x86	N-DIM [3] PLC GPIB Other network-based devices	Soft IOC	24
Linux x86	CAMAC	Embedded IOC	6
Linux f3RP61 [4]	PLC	Embedded IOC	14
vxWorks	NIO [5]	Embedded IOC	7

PV MANAGEMENT SYSTEM

IOC and Shared Storage

In the RIBF control system, common EPICS programs (EPICS-base, application programs, runtime database, and additional extensions programs) are stored in the network attached storage (NAS), and they are shared by all EPICS IOCs using the network file system or file transfer protocol. The system diagram is shown in Fig. 1. The NAS, manufactured by NetApp, provides a centralized system with a high-availability. On the other

#a-uchi@riken.jp

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

THE POWER SUPPLY CONTROL SYSTEM OF CSR

W. Zhang, S.An, S.Z.Gou, K.W.Gu, Y.P.Wang ,P. Li, M.Yue
IMP, LAN Zhou 730000, P.R. China

Abstract

This article gives a brief description of the power supply control system for Cooler Storage Ring (CSR). It introduces in detail mainly of the control system architecture, hardware and software. We use standard distributed control system (DCS) architecture. The software is the standard three-layer structure. OPI layer realizes data generation and monitoring. The intermediate layer is a data processing and transmission. Device control layer performs data output of the power supply. We use ARM + DSP controller designed by ourselves for controlling the power supply output. At the same time, we have adopted the FPGA controller designed for timing for power supply control in order to meet the requirements of accelerator synchronizes the output of the power supply.

INTRODUCTION

HIRFL-CSR, a new ion Cooler-Storage-Ring (CSR) project, is the post-acceleration system of the Heavy Ion Research Facility in Lanzhou (HIRFL). HIRFL-CSR is a multi-purpose CSR system that consists of a main ring (CSRm), an experimental ring (CSRe), and a radioactive beam line (RIBLL II) to connect the two rings. Figure 1 show an Overall Layout of HIRFL-CSR [1]. The two existing cyclotrons SFC (K = 69) and SSC (K = 450) of the HIRFL will be used as its injector system. The heavy ion beams with the energy range of 8–30 MeV/u from the HIRFL will be accumulated, cooled and accelerated to the high-energy range of 100–400 MeV/u in the main ring, and then extracted fast to produce RIB or highly charged heavy ions. The secondary beams (RIB or highly charged heavy ions) will be accepted and stored by the experimental ring for many internal-target experiments or high-precision spectroscopy with beam cooling. On the other hand, the beams with the energy range of 100–900 MeV/u will also be extracted from CSRm by using slow extraction or fast extraction for many external-target experiments.

Two electron coolers located in the long-straight sections of CSRm and CSRe, respectively, will be used for the beam accumulation and cooling. One internal target in the long-straight section of CSRe will be used for nuclear physics and highly charged state atomic physics, and many external targets of CSRm will be used for nuclear physics, cancer therapy study and other researches.

CSR is a double ring system. In every operation cycle, the stable-nucleus beams from the injectors are accumulated, cooled and accelerated in the main ring (CSRm), then extracted fast to produce RIB or highly charged ions. The experimental ring (CSRe) can obtain the secondary beams once for every operation cycle. The accumulation duration of CSRm is about 10 s. considering the ramping rate of magnetic field in the dipole magnets to be 0.1–0.4 T/s, the acceleration time of CSRm will be nearly 3 s. Thus, the operation cycle is about 17 s [2].

In CSRe, two operation modes will be adopted. One is the storage mode used for internaltarget experiments or high-precision spectroscopy with electron cooling. Another one is the deceleration-storage mode used for atomic-physics experiments.

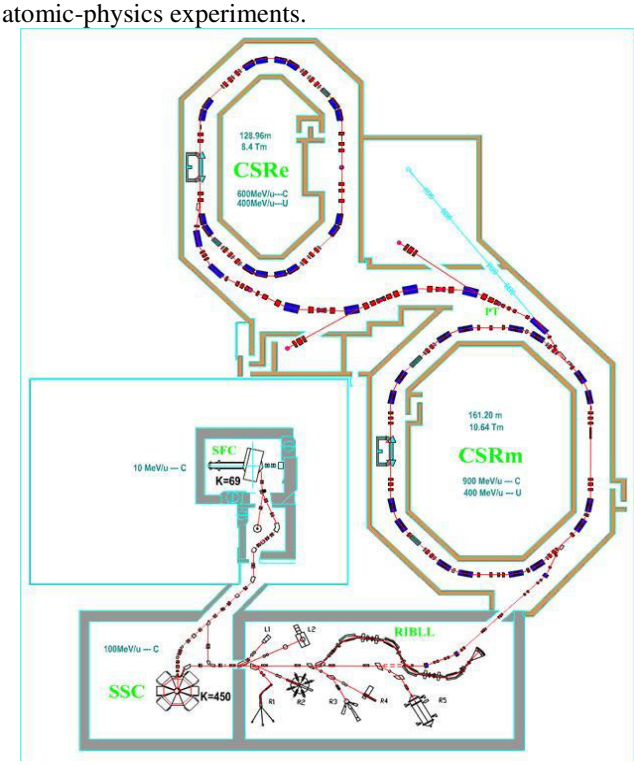


Figure 1: Overall Layout of HIRFL-CSR.

SYSTEM INTRODUCTION

The control system is based on the Ethernet star topology. It takes Ethernet as the transmission medium to connect each part, mainly including: database system, the synchronization server, the front end server, the I/O part (ARM controller and the DSP processor) and the communication module. Database system and synchronization server is the information and control centre of the whole system. There is an independent Ethernet connection between the front end of the server and the I/O component, and a front-end server manages multiple I/O components. The I/O part directly controls the device object, and an I/O part controls one or four device objects. The structure of the power control system is shown in Figure 2. In the system, power control is accomplished by three kinds of data files: process data, event table data and command table data.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

DATA CATEGORIZATION AND STORAGE STRATEGIES AT RHIC*

S.Binello†, K.Brown, T.D'Ottavio, J.Laster, R.Katz, J.Morris, J.Piacentino
Collider-Accelerator Department, BNL,Upton NY, USA

Abstract

This past year the Controls group within the Collider Accelerator Department at Brookhaven National Laboratory replaced the Network Attached Storage (NAS) system that is used to store software and data critical to the operation of the accelerators. The NAS also serves as the initial repository for all logged data. This purchase was used as an opportunity to categorize the data we store, and review and evaluate our storage strategies. This was done in the context of an existing policy that places no explicit limits on the amount of data that users can log, no limits on the amount of time that the data is retained at its original resolution, and that requires all logged data be available in real-time. This paper will describe how the data was categorized, and the various storage strategies used for each category.

INTRODUCTION

The NAS that had served as the main repository of data and software critical to running the accelerators, as well as the initial repository of logged data since 2008 was declared End Of Life by the vendor in 2014. As such, we needed to decide whether to continue using the system and obtain support from a third party, or simply replace the system. As this NAS had other limitations: insufficient storage, outdated small expensive 300GB drives and 1GB network interfaces, along with a hefty $15K annual maintenance fee for a mere 25TB of storage, the decision was made to replace the NAS.

DATA CATEGORIZATION

To facilitate and organize the analysis of requirements for the new system we categorized our data into three tiers: critical operational data, auxiliary operational data and historical/temporary data. Each tier was then assigned a level of required reliability.

Tier 1 – Critical Operational Data

This is data that is needed to run the accelerators. It consists of program executables, basic configuration information, Front End Computer (FEC) boot areas, archives of recent settings, and critical group home directories.

This data should always be available. A complete High Availability(HA) solution is required. That is, the storage device should have redundant components at all levels. In the event of multiple failures which makes this data unavailable, a backup with data as current as possible should be made available. A switch to the backup system should take less than 4 hours.

*Work performed under Contract Number DE-AC02-98CH0886 with the auspices of the US Department of Energy
† sev@bnl.gov

Tier 2 – Auxiliary Operational Data

This data may be important for some operational tasks, but is not critically needed for basic machine operation, for example, logged data from the current run, archived settings from past runs (and possibly archived settings from earlier in this run).

A total HA solution is desired, but is not absolutely necessary. Minimally, RAID should be used to protect from disk failures. If primary storage for this data is not a total HA solution, a backup version of the same data should be maintained on an alternate disk storage device. Data on the backup should be no more than three days old. A switch to the backup should take less than 4 hours.

Tier 3 – Historical or Temporary Data

This data is mostly composed of logged data from past runs, along with any data that may only need to be stored temporarily.

Some high availability features are desired. Minimally, RAID should be used to protect from disk failures. In the event of some other failure, restoration from tape would be necessary. Restoration of most types of data could be done in less than an hour but could take several days.

DATA STORAGE STRATEGIES

Once the data was classified attention focused no how best to support the requirements for each data tier.

Tier 1 – Storage For Critical Data

From the onset, it was understood that a High Availability (HA) NAS was the required solution for critical operational data. The focus was then on how to increase the availability of this critical data in the event of a failure of the primary NAS, and how to provide the most up-to date data possible. The Disaster Recovery (DR) system for the NAS that was to be replaced, was a low cost Linux storage server with internal SATA drives RAID'ed using 3-ware controllers (this is a configuration that is also use to support tier 3 data). Critical data was replicated from the NAS to the DR system using the Linux rsync utility. This pairing proved to have some flaws. The DR system was not identical to the NAS, and as the NAS was asked to store and perform more and more, over time it eventually outpaced the capabilities of the DR system. Additionally, we found that the rsyncs were also negatively impacting the performance of the primary NAS.

It was decided that in order to prevent this type of divergence in capabilities in the future, and in an attempt to limit the impact of data replication between the primary and DR unit, an identical NAS was needed for disaster

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

DATA LIFECYCLE IN LARGE EXPERIMENTAL PHYSICS FACILITIES: THE APPROACH OF THE SYNCHROTRON ELETTRA AND THE FREE ELECTRON LASER FERMI

Fulvio Billè#, Roberto Borghes, Francesco Brun, Valentina Chenda, Alessio Curri, Venicio Duic, Daniele Favretto, Georgios Kourousias, Marco Lonza, Milan Prica, Roberto Pugliese, Martin Scarcia, Michele Turcinovich, Elettra Sincrotrone Trieste, Basovizza, Trieste, Italy

Abstract

Often the producers of Big Data face the emerging problem of Data Deluge. The literature includes experiences and proposals for dealing with this problem. Nevertheless experimental facilities, such as synchrotrons and free electron lasers, may have additional and specialised requirements. This is partially due to the plethora of cutting-edge custom instrumentation but also to the necessity for a complex system that manages the access to the facility for thousands of scientists. A complete data lifecycle describes the seamless path that joins distinct IT tasks such as experiment proposal management, user accounts, beamline software, data acquisition, reduction and analysis, archiving, cataloguing, and remote access. This short paper presents the data lifecycle of the synchrotron Elettra and the free electron laser FERMI. It attempts to go beyond the overgeneralised binary distinction of information between data and metadata since in practice there are many heterogeneous but related data sources. With the focus on the broad concept of data access, the Virtual Unified Office is presented. It is a core element involved in scientific proposal management, user information database, scientific data oversight, and remote access. Eventually recent choices of advanced beamline software are here discussed as well as approaches to distributed control systems. The latter holds the key role to data and metadata acquisition but it also requires integration with the rest of the system components in order to facilitate additional services such as cataloguing, archiving and remote access. With the experience of the past years and the forecasting of future requirements, the expected system upgrades are also here outlined. The scope of this paper is to disseminate the current status of a complete data lifecycle, discuss key issues and hints for future directions. This is potentially useful for similar facilities especially for those under construction or upgrade.

INTRODUCTION

Elettra and FERMI

The research centre Elettra consists of two advanced light sources, the electron storage ring Elettra and the free-electron laser (FEL) FERMI hosting 32 experimental stations.

Elettra is a third-generation synchrotron serving the scientific and industrial community since 1993. It is involved in a broad spectrum of research in physics, chemistry, biology, life sciences, environmental science, medicine, forensic science, and cultural heritage. Currently 28 beamlines, including a storage-ring free-electron laser, utilize the radiation generated by the Elettra source. All of the most important x-ray based techniques in the areas of spectroscopy, spectromicroscopy, diffraction, scattering and lithography are present, together with facilities for infrared microscopy and spectroscopy, ultraviolet inelastic scattering, and band mapping. In 2010 the third-generation electron storage ring source Elettra was upgraded to operate in top-up mode. It is the only third-generation synchrotron radiation source in the world that operates routinely at two different electron energies; i) 2.0 GeV for enhanced extended ultraviolet performance and spectroscopic applications, and ii) 2.4 GeV for enhanced x-ray emission and diffraction applications. A plan for a substantial upgrade named Elettra 2.0 is in motion [1].

FERMI is a new seeded free electron laser (FEL) facility in operation next to the third-generation synchrotron radiation facility. Unique among the FEL sources currently operating in the ultraviolet and soft x-ray range worldwide, FERMI has been developed to provide fully coherent ultrashort 10-100 femtosecond pulses with a peak brightness ten billion times higher than that made available by third-generation light sources. It is a single-pass FEL that covers the wavelength range from 100 nm to 4 nm in the first harmonic. It comprises two separate coherent radiation sources, FEL-1 and FEL-2, that are being brought online sequentially. FEL-1 operates in the wavelength range between 100 and 20 nm via a single cascade harmonic generation, while the latest FEL-2 is designed to operate at shorter wavelengths (20-4 nm) via a double cascade configuration.

Data Deluge and Discontinuity

The era of Big Data has raised relevant problems like that of Data Deluge [2]. Elettra and similar facilities can be seen as information generators. The processes that generate such information flows are rather specialised since the facilities themselves have particular workflows. Such workflows include proposal submission, evaluation, beamline access, execution of experiments with custom instrumentation, data analysis, dissemination of scientific results etc. At this stage the issue of Data Deluge is studied in the focused context of a synchrotron facility.

#fulvio.bille@elettra.eu

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

A FLEXIBLE SYSTEM FOR END-USER DATA VISUALISATION, ANALYSIS PROTOTYPING AND EXPERIMENT LOGBOOK

Roberto Borghes*, Valentina Chenda, Georgios Kourousias, Marco Lonza, Milan Prica, Martin Scarcia, Elettra-Sincrotrone Trieste S.C.p.A, Basovizza, Italy

Abstract

Experimental facilities like synchrotrons and free electron lasers, often aim at well defined data workflows tightly integrated with their control systems. Still such facilities are also service providers to visiting scientists. The hosted researchers often have requirements different than those present in the established processes. The most evident needs are those for i) flexible experimental data visualisation, ii) rapid prototyping of analysis methods, and iii) electronic logbook services. This paper reports on the development of a software system, collectively referred to as DonkiTools, that aims at satisfying the aforementioned needs for the synchrotron Elettra and the free electron laser FERMI. The design strategy is outlined and includes topics regarding: dynamic data visualisation, Python scripting of analysis methods, integration with the TANGO distributed control system, electronic logbook with automated metadata reporting, usability, customization, and extensibility. Finally a use case presents a full deployment of the system, integrated with the FermiDAQ data collection system, in the free electron laser beamline EIS-TIMEX.

INTRODUCTION

At the present moment there are 32 experimental stations operating at the synchrotron Elettra and the free electron laser (FEL) FERMI [1]. The data acquisition and control software of each end-station is almost unique, being highly focused on the specific technique applied. Nevertheless, for the user-oriented applications it is not unusual to meet the requirements that are common to a number of end-stations. For those, the most efficient solution would be a product with a high level of customization reusable in a multitude of situations. As a general idea, such an approach may be compared to a drill tool: it has been designed to perform a specific task but it needs to mount accessories for the alternative uses like drilling in metal, wood or concrete.

The Scientific Computing team had the chance to apply this strategy in different situations, building a collection of software tools called the DonkiTools. Any of the DonkiTools aims at becoming a standard in its field of application at Elettra and FERMI. The name of the suite recalls that of a donkey, a robust and versatile beast of burden, so highly appreciated in the Mediterraen region.

* roberto.borghes@elettra.eu

DONKITOOLS OVERVIEW

Each application of the DonkiTools suite (Fig. 1) originated to solve a specific requirement, but has been designed from the start using a plug-in based architecture that makes it flexible and reusable. The plugin-based extensibility allows for constant evolution of the DonkiTools. It is possible to add a new functionality without the hazard of destabilizing others or the core behaviour.

As a programming language, Python was the natural choice as it is supported by TANGO and offers a high degree of dynamicity and a large set of scientific and visualization libraries.

PyQt4 [2] has been adopted as the GUI development framework. PyQt ensures that the DonkiTools will run on all platforms supported by the Qt, including Windows, MacOS/X and Linux. The plug-ins that can extend each application are basically PyQt Widgets included in the main panel at runtime.

In the next sections, we will present a short overview of three successful applications used by the staff and visiting researchers at Elettra and FERMI:

- DonkiLOG: an extensible electronic logbook for experiments with automated metadata reporting.
- DonkiCCD: a flexible visualizer of scientific images that may be customized and integrated to the end-station control system.
- DonkiPY: a tool for rapid prototyping of data collection and analysis scripts.



Figure 1: DonkiTools icons.

DONKILOG

The first release of the DonkiLOG (Fig. 2) was developed for the needs of scientists working at the free electron laser FERMI. Each experimental shift at FERMI consists of several scan operations. Scans are often

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

MONITORING MIXED-LANGUAGE APPLICATIONS WITH ELASTICSEARCH, LOGSTASH AND KIBANA (ELK)

A. De Dios Fuente, O. O. Andreassen, C. Charrondière, CERN, Geneva, Switzerland

Abstract

Application logging and system diagnostics is nothing new. Ever since we had the first computers scientists and engineers have been storing information about their systems, making it easier to understand what is going on and, in case of failures, what went wrong. Unfortunately there are as many different standards as there are file formats, storage types, locations, operating systems, etc. Recent development in web technology and storage has made it much simpler to gather all the different information in one place and dynamically adapt the display. With the introduction of Logstash with Elasticsearch as a backend, we store, index and query data, making it possible to display and manipulate data in whatever form one wishes. With Kibana as a generic and modern web interface on top, the information can be adapted at will. In this paper we will show how we can process almost any type of structured or unstructured data source. We will also show how data can be visualised and customised on a per user basis and how the system scales when the data volume grows.

INTRODUCTION

At CERN, as in any other large organization, lots of data are generated and stored at a rate that makes it difficult for humans to analyse them. In addition, when considering the many different file formats, storage types, and physical locations, treating the data manually can become an unfeasible and overwhelming task. Manual treatment would go from data analysis to bug or error tracking. Having a tool that helps engineers understand the insights of the data stored, generate reports and gather statistics from them in a faster and easier way, is often the key to prevent future problems and where the challenge lies.

Moreover, with any kind of system, the logfile is the first place to look for clues when the system behaves in an unexpected manner. In our case, we have a variety of applications written in LabVIEW, C++ and Java, running on the most popular operating systems (Linux, Windows and OS X) and applications that make use of our LabVIEW Rapid Application Development Environment (RADE) [1], which has part of the services distributed over multiple servers and which is accessible from more than ~100 users at CERN. In such systems, most of the logs have different or weak structures. This led us to investigate how to unify all the different data formats and sources, and how to store all the information in a common place in order to have a tool that monitors and keeps track of all the online data easily and effectively in real time. However, since the needs of the users are different according to the application in hands, the way of getting statistics and reports should be customizable without web-development knowledge or specialised skills.

According to the motivation explained above, a conscientious study was made, in which several tools were identified, studied and tested.

STUDY OF TOOLS

One of the first things to consider was the data format. There is a large variety of sources; LabVIEW applications running on CompactRIO and PXI targets, Apache Tomcat servers, Java services, C++ applications and extensions, where each of them have different formats and purposes. In computing, syslog is a widely used standard for message logging [2] so this was the starting point to define the main format. In addition, the main requirements for the desired system are listed below:

- Support multiple log formats but mainly syslog.
- Support different communications and network protocols.
- Centralised data messages.
- Have a web-viewer to analyse the logs.
- Be able to get statistics of the stored data.
- Be fully compatible with Linux Environment.
- Be easy to use, install and configure.
- Be scalable if data grows over time.

Analysis of Logging Tools

Considering the advances in data mining and the fact that the analysis of “Big Data” [3] is not a new field, we decided to look for an already existing tool.

After searching and evaluating some promising tools found in the market, Logstash and Fluentd [4] were selected for a deeper study and test. Since these tools are evolving and have been redesigned several times due to their increasing popularity, some characteristics and features might have changed after our initial test. The versions used were Logstash 1.3, Fluentd 1.1, Elasticsearch 0.90 and Kibana 3.

Logstash and Fluentd are aimed to unify and manage data collection for better use and understanding. Both are free, open source and based in plugin models to extend functionality. The two of them can be combined with the popular Elasticsearch as the backend data store and Kibana as a front-end reporting tool to complete all the requirements.

Logstash is based on inputs, filters, codecs and output plugins where the inputs are the sources of the data, the filters are processing actions on the data under certain conditions, the codecs change the data representation and finally the outputs are the destinations where the data is sent. Fluentd has the same behaviour for the inputs and outputs but internally all the data are converted to JSON

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

SCALABLE WEB BROADCASTING FOR HISTORICAL INDUSTRIAL CONTROLS DATA

B. Copy, O. Andreassen, P. Gayet, M. Labrenz, H. Milcent, F. Piccinelli
CERN, Geneva, Switzerland

Abstract

With the widespread adoption of asynchronous web communication, thanks to WebSockets and WebRTC [1], it has now become possible to distribute industrial controls data (such as coming from devices or SCADA software) in a scalable and event-based manner to a large number of web clients [2] in the form of rich, interactive visualizations. There is yet, however, no simple, secure and performant way to query large amounts of aggregated historical data.

This paper presents an implementation of such an architecture, which is able to make massive quantities of pre-indexed historical data stored in ElasticSearch available to a large number of web-based consumers through asynchronous web protocols. It also presents a simple, Opensocial-based [3] dashboard architecture that allows users to configure and organize rich data visualizations (based on Highcharts Javascript libraries) and create navigation flows in a responsive, mobile-friendly user interface.

Such techniques are used at CERN to display interactive reports about the status of the LHC infrastructure (*e.g.*, Vacuum and Cryogenics installations) and give access to fine-grained historical data (such as stored in the LHC Logging database [4]) in a matter of seconds.

THE CERN ENGINEERING DASHBOARD

The CERN Engineering Dashboard is primarily a data streaming facility, giving secure and scalable web-based access to live, critical data coming directly from the CERN Accelerator infrastructure.

The CERN Engineering Dashboard also promotes the usage of modern, standards-based data visualizations, leveraging the best open web standards (*e.g.*, CSS 3, HTML 5, SVG), thereby ensuring that the solution can cope with technological debt and clearly separate the data broadcasting mechanism from the way data is rendered.

Figure 1 below gives an overview of the Broadcasting Dashboard architecture, allowing the scalable distribution of both live and historical data and the rendering of said data through modular and reusable visualizations.

STREAMING HISTORICAL DATA

While our previous publication on the topic of web broadcasting (MOPPC145 [1]) focused on the ability to stream live data, the CERN Engineering Dashboard project has aimed for the past two years to provide the same functionality for historical data, introducing in the process a whole new set of challenges:

- Instantly retrieving and serving large amounts of data, instead of atomic values, while still being able to serve in the vicinity of one thousand web clients concurrently.
- Aggregating data on the fly so that mobile clients are not submerged with unnecessary information that they will, in any case, be unable to display on a small viewport.
- Not making any assumptions on the data at hand, yet still providing efficient indexing and filtering capabilities.

Over the course of this project, it soon became obvious that traditional data storage techniques, such as relational databases, were insufficient:

- Connections to relational databases must usually be pooled (about a dozen concurrent connections are tolerated on the central CERN Oracle database); serving the needs of thousands of clients concurrently was not possible.
- Relational data access times impose dedicated connection drivers and too much latency, and transmitting this data over the web imposes yet more delay and processing cost in data format translation.

A new generation of data storage solutions, identified under the umbrella term “NoSQL,” has emerged in the past few years. NoSQL solutions (thus coined for their breaking free from traditional SQL relational data storage constraints) rely on asynchronous, massively parallel processing and cloud clustering techniques in order to address the aforementioned requirements.

NOSQL DATA STORAGE SOLUTIONS

There are several NoSQL data storage solutions on the market at present. Our goal was to identify a solution that:

- was web-friendly, could natively communicate via HTTP, and support a query language as expressive and feature-capable as SQL;
- could operate in cluster mode, so as to increase availability and support data partitioning;
- could integrate with online processing frameworks such as Apache Storm;

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

METADATASTORE: A PRIMARY DATASTORE FOR NSLS-2 BEAMLINES

A. Arkilic, L. Dalesio, D. Allan, W.K. Lewis, T. A. Caswell, NSLSII, Upton, NY, USA

Abstract

The beamlines at NSLS-II are among the highest instrumented, and controlled of any worldwide. Each beamline can produce unstructured data sets in various formats. This data should be made available for data analysis and processing for beamline scientists and users. Various data flow systems are in place in numerous synchrotrons, however these are very domain specific and cannot handle such unstructured data. We have developed a data flow service, metadatastore that manages experimental data in NSLS-II beamlines. This service enables data analysis and visualization clients to access this service either directly or via databroker API in a consistent and partition tolerant fashion, providing a reliable and easy to use interface to our state-of-the-art beamlines.

INTRODUCTION

Historically, the experimental data in beamline experiments are either hosted in purpose-specific relational database management systems or text files. Relational database systems generally provide robust and high performance solutions for most experiments. However, they can only be adopted by certain experiments since two experiments in the same experimental floor usually have very different data formats and indexing requirements. Using text and/or hierarchical file systems such as HDF5 addresses some of these needs. File based data storage has served the legacy system in NSLS-I for almost 3 decades but the dramatic increase in data rates of NSLS2 beamlines made file I/O a massive bottleneck. In order to address these issues and provide NSLS2 beamlines with a uniform data storage solution that is applicable to all our beamlines, we have developed metadatastore service using NoSQL backend MongoDB [1].

ARCHITECTURE

metadatastore is the primary source of storage for scans, sweeps, etc. metadatastore is implemented as a RESTful web service on Tornado with a NoSQL MongoDB backend. Since it is document based, MongoDB does not have the traditional notion of tables. Each entry to the database is a document that is analogous to an "entry" in relational database terms. MongoDB collections serve as the containers of documents, providing its clients with means to meaningfully group their documents [2]. The metadatastore backend consists of 4 major collections: RunStart, EventDescriptor, Event, RunStop. One can think of RunStart and RunStop as the head and tail of a series of events that happen throughout an experiment (Figure 1). This approach eliminates the need to update the header of the experiment, which eliminates the problems that could occur due to lack of transactions in MongoDB. EventDescriptor(s) contain information regarding the data acquisition Event(s). In other words, they contain information about what are the

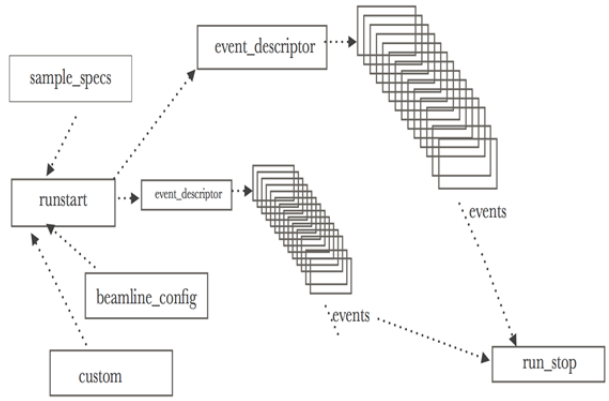


Figure 1: metadatastore Collection Relationship Overview.

"data_keys" that are being captured, what their shapes and data types are. This is immensely useful for data analysis toolkit in NSLS2 as this sort of header information is a major requirement to set up the data analysis environment prior to the arrival of actual data payload. In addition to the advantages provided so far, because of immense flexibility and performance of MongoDB, these RunStart, RunStop, and EventDescriptor documents support addition of any name-value pair or object in any hierarchy. Since Event(s) contain the actual data payload, they are defined as immutable documents. This semi-structured design pattern allows NSLS-II middle-layer to tackle many challenges including asynchronous scans, custom hardware brought from an external source, and varying data formats among experiments. For instance, a coherent X-ray diffraction and protein crystallography beamline utilize metadatastore in order to save their experiments despite their very different use-cases.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

FILESTORE: A FILE MANAGEMENT TOOL FOR NSLS-II BEAMLINES

A. Arkilic, L. Dalesio, D. Chabot, W.K. Lewis, D. Allan, T. A. Caswell
Brookhaven National Laboratory, NSLSII, Upton, NY, USA

Abstract

NSLS-II beamlines can generate 72,000 data sets per day resulting in over 2 M data sets in one year. The large amount of data files generated by our beamlines poses a massive file management challenge. In response to this challenge, we have developed filestore, as means to provide users with an interface to stored data. By leveraging features of Python and MongoDB, filestore can store information regarding the location of a file, access and open the file, retrieve a given piece of data in that file, and provide users with a token, a unique identifier allowing them to retrieve each piece of data. filestore does not interfere with the file source or the storage method and supports any file format, making data within files available for NSLS-II data analysis environment.

INTRODUCTION

Current state-of-the-art detectors are capable of producing megapixel images with frame-rates in the kilohertz range which may result in peak data rates of up to 800 Mb.s-1. As a result, the volume of data associated with an experiment may now exceed several tens of terabytes so simply moving data from one location to another is prohibitive. As such, it may not be possible for external users to take a copy of all the data on portable media or to easily transfer over the Internet. Thus it requires high-performance data storage systems and advanced processing capabilities to allow users to access their data both during their experiment at NSLS-II and to continue the analysis following the NSLS-II visit. The middle-layer architecture of NSLS-II beamlines provides such capabilities using a blend of NoSQL databases and distributed file systems [2]. **fileStore** is the component of the middle-layer that keeps track of experimental files. These experimental files are generated via EPICS areaDetector module or any custom camera. filestore's NoSQL mongo backend allows users to save any sort of information about their images, while providing links to metadatastore runs[1]. Filestore provides clients with high-performance data mining techniques that were not possible via legacy applications in place.

ARCHITECTURE

filestore provides information regarding the location of the file within the General Parallel File System (GPFS). MongoDB serves as the back-end to filestore due to its flexibility and performance. Just like metadatastore, filestore has a set of required standard fields and allows users/data acquisition scripts to add any field in any hierarchy, making it a primary storage environment for image headers [1].

```
In [1]: from filestore.api import insert_resource, insert_datum

        resource_id = insert_resource('csv', 'example.csv')
```

Figure 1: filestore Python API Resource Insert Example.

Filestore's capabilities far exceed providing URL information. It also serves as a platform for performing data file formatting (via handlers) and data regression (via its complex querying capabilities). Experimental files can be stored completely independent of filestore's acknowledgement. Since filestore does not need to be aware of the file at time of creation, it is possible to migrate data files from legacy systems into the modern middle-layer service framework of NSLS-II beamlines (see Fig. 1).

```
In [5]: insert_datum(resource_id, 'some_id1', {'line_no': 1})
Out[5]: <Datum: Datum object>

In [6]: insert_datum(resource_id, 'some_id2', {'line_no': 2})
Out[6]: <Datum: Datum object>

In [7]: insert_datum(resource_id, 'some_id3', {'line_no': 3})
Out[7]: <Datum: Datum object>

In [8]: insert_datum(resource_id, 'some_id4', {'line_no': 4})
Out[8]: <Datum: Datum object>

In [9]: insert_datum(resource_id, 'some_id5', {'line_no': 5})
Out[9]: <Datum: Datum object>
```

```
In [10]: from filestore.api import register_handler

In [11]: register_handler('csv', CSVLineHandler)
```

Finally, we are ready to retrieve that data. All we need is the unique ID.

```
In [12]: from filestore.api import retrieve

In [13]: retrieve('some_id2')
Out[13]: 'b\n'
```

Figure 2: Datum and Handler example.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

LARGE GRAPH VISUALIZATION OF MILLIONS OF CONNECTIONS IN THE CERN CONTROL SYSTEM NETWORK TRAFFIC: ANALYSIS AND DESIGN OF ROUTING AND FIREWALL RULES WITH A NEW APPROACH

Luigi Gallerani, CERN, Geneva, Switzerland

Abstract

The CERN Technical Network (TN) TN was intended to be a network for accelerator and infrastructure operations. However, today, more than 60 million IP packets are routed every hour between the General Purpose Network (GPN) and the TN, involving more than 6000 different hosts. In order to improve the security of the accelerator control system, it is fundamental to understand the network traffic between the two networks and to define new appropriate routing and firewall rules without impacting operations. The complexity and huge size of the infrastructure and the number of protocols and services involved, have discouraged for years any attempt to understand and control the network traffic between the GPN and the TN. In this paper, we show a new way to solve the problem graphically. Combining the network traffic analysis with the use of large graph visualization algorithms we produced usable 2D large color topology maps of the network identifying the inter-relations of the control system machines and services, in a detail and clarity, not seen before.

INTRODUCTION

Improving the security between the GPN and the TN, where accelerator systems and control systems run, is a priority for the infrastructure experts in the CERN Beams Department, Control Group, and the Computer Security team.

The separation between the GPN and the TN is defined by routing rules controlled by the Network Operations database (NetOps or LanDB). Only GPN hosts that are TN-TRUSTED are able to connect to the TN, and only TN hosts that are GPN-EXPOSED can communicate with the GPN. From the CERN NetOps records, routing tables are generated and loaded in to the routers.

We have today more than 1000 machines that are in the TRUSTED or EXPOSED lists with a justified operational reason. Around 6000 machines are involved in the communications between the two networks. The TRUSTED – EXPOSED list mechanism is clearly too weak compared to the modern network protection standards. The goal is to have in place well defined firewall and routing rules and have full control on the traffic between TN and GPN. To achieve that on a large control system in operation, it is fundamental to and understand all the relations and dependencies between the system involved.

A good starting point is the record of the real traffic on routers.

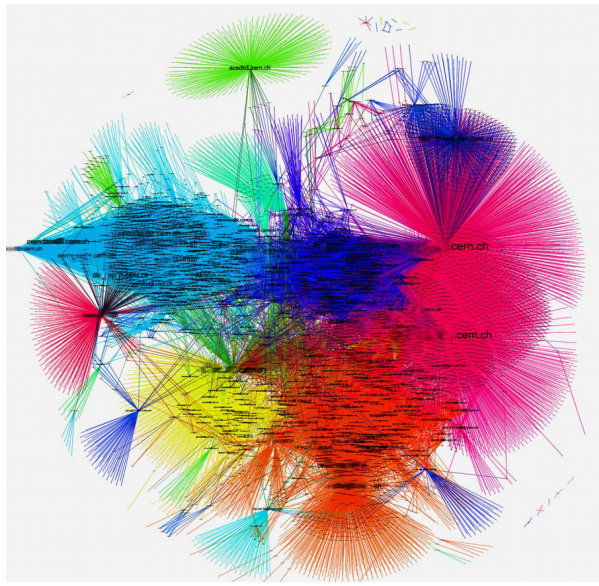


Figure 1: One hour of UDP-TCP traffic between CERN Technical Network and General Purpose Network plotted and clustered using Fruchterman Reingold algorithm.

RECORDING THE TRAFFIC

A large set of network restrictions rules can be easily defined by the knowledge of the system. However, this is not sufficient. With hundreds of different interconnected systems running on the TN and GPN, it is almost impossible to predict all the relations and dependencies between hosts and ports by the knowledge of the system functionality. The risk of compromising the connectivity of a running system because of a too strict or wrong firewall rule is what the Controls group wants to avoid most. Routers between the TN and GPN can record the traffic live producing as output text files with the following fields

Timestamp	Protocol	Packets count	Size bytes
Source host	Source port	Dest host	Dest port

DATA PROCESSING

Classical Approach by Statistics

In 1 hour, 60 million connections are recorded. 1.4 billion lines per day are the source of the analysis. Statistics, filtering, grouping and sorting are the obvious approaches for useful information extraction from a large dataset. For example it is easy to identify by queries the most connected hosts, the most used port between hosts of name N, the top 10 interconnected hosts and so on.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

TOWARDS A SECOND GENERATION DATA ANALYSIS FRAMEWORK FOR LHC TRANSIENT DATA RECORDING

S. Boychenko, C. Aguilera-Padilla, M. Dragu, M.A. Galilee, J.C. Garnier, M. Koza, K. Krol, R. Orlandi, M.C. Poeschl, T.M. Ribeiro, K. Stamos, M. Zerlauth CERN, Geneva, Switzerland
M.Z. Rela CISUC, University of Coimbra, Portugal

Abstract

During the last two years, CERNs Large Hadron Collider (LHC) and most of its equipment systems were upgraded to collide particles at an energy level twice higher compared to the first operational period between 2010 and 2013. System upgrades and the increased machine energy represent new challenges for the analysis of transient data recordings, which have to be both dependable and fast. With the LHC having operated for many years already, statistical and trend analysis across the collected data sets is a growing requirement, highlighting several constraints and limitations imposed by the current software and data storage ecosystem. Based on several analysis use-cases, this paper highlights the most important aspects and ideas towards an improved, second generation data analysis framework to serve a large variety of equipment experts and operation crews in their daily work.

INTRODUCTION

The consolidation and upgrade activities during the recent long shutdown has allowed almost doubling the collision beam energy to 13 TeV. This implies as well a much increased damage potential in case of the serious failure in comparison to previous runs. As the rest of the machine, the systems which protect the LHC have equally undergone major improvements based on the experience collected during the first operational run between 2010 and 2013. Besides the increased number of machine protection devices installed in the machine, the amount of data which they produce also increased. The software which is used for analysis of accelerator diagnostic data requires a major upgrade to keep up with the LHC expansion and the unprecedented amounts of data to be processed [1].

Based on experience of the previous hardware commissioning we have identified the main shortcomings with the current data analysis approach. Combined with the use cases collected over several operational years and the feedback provided by operators and hardware experts, we were able to define the main goals and requirements for the next generation accelerator data analysis framework. In order to solve the identified problems, an extensive study was performed, followed by a comparison between existing data storage and processing tools. Although, the integration of modern data analysis frameworks might improve the situation with respect to the currently implemented approach, we believe that specificities of collected data and environment itself pose a greater challenge. Based on the initial research, a new workload-driven approach for organizing the data is proposed, aiming to serve efficiently heterogeneous workloads. The strengths and the weaknesses of the proposed solution are being analysed, as well as we discuss the details of the integration of the proposed technique with currently deployed accelerator software stack. Finally, the goals for the future research and development of the new analysis framework are defined.

The remaining document is organized as follows: next section provides a short overview on the current analysis framework shortcomings and desired requirements for the new solution. The following section describes the novel approach which is believed to allow building the system capable of satisfying the majority of the use cases. In future work we outline the goals for future research and development. Finally in the last section a summary on the main conclusions is presented.

CURRENT ANALYSIS FRAMEWORK

During the last few years, significant efforts were made to initiate the process of LHC transient data analysis centralization. The operators and equipment experts were previously relying on multiple service-specific graphical tools to study the underlying data. The process required significant effort and concentration, especially during accelerator commissioning when thousands of tests are executed in order to validate the behaviour of control systems. Furthermore, the data correlation between different storages was either performed manually or required the development of custom modules, which would take care of the required merging process. A recent attempt on automatizing the validation process for LHC protection systems resulted in the organization and systematization of the analysis process. The results were consolidated into the AccTesting framework [2] which has proven its efficiency during the last LHC hardware commissioning phase.

Infrastructure

There are two major sources of LHC diagnostic data: The Post Mortem (PM) framework [3] and CERNs Accelerator Logging Service (CALS) [4]. Despite acquiring the similar data from a given LHC device, the two services have fundamentally different use cases and properties. The main differences between the two frameworks is the resolution of collected data. The Post Mortem system only collects data around interesting events, by retrieving all the entries from the internal buffer of the monitored devices.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

SMOOTH MIGRATION OF CERN POST MORTEM SERVICE TO A HORIZONTALLY SCALABLE SERVICE

C. Aguilera-Padilla, S. Boychenko, M.Dragu, M.A. Galilee, J.C. Garnier, M. Koza, K. Krol, R. Orlandi, M.C. Poeschl, T.M. Ribeiro, M. Zerlauth, CERN, Geneva, Switzerland

Abstract

The Post Mortem service for the CERN accelerator complex stores and analyses transient data recordings of various equipment systems following certain events, like a beam dump or magnet quenches. The main purpose of this framework is to provide fast and reliable diagnostic to the equipment experts and operation crews to help them decide whether accelerator operation can continue safely or whether an intervention is required. While the Post Mortem System was initially designed to serve the CERN Large Hadron Collider (LHC), its scope has been rapidly extended to include as well External Post Operational Checks and Injection Quality Checks in the LHC and its injector complex. These new use cases impose more stringent time-constraints on the storage and analysis of data, calling for a migration of the system towards better scalability in terms of storage capacity as well as I/O throughput. This paper presents an overview of the current service, the ongoing investigations and plans towards a scalable data storage solution and API, as well as the proposed strategy to ensure an entirely smooth transition for the current Post Mortem users.

INTRODUCTION

The Post Mortem (PM [1]) service has been providing data collection, storage and analysis of LHC event data since 2008. The PM system usage grew significantly from the first use cases covering the understanding of the conditions in which the machine aborted its operations: fault on an equipment, powering event, etc. It currently covers the Injection Quality Checks (IQC [2]) from the SPS to the LHC and the Extraction Post Operational Checks (XPOC [3]) to verify that the beam extraction from the LHC was correctly executed. A new use case was just implemented to analyze every single cycle of the SPS. A possible integration of LINAC4 is under study. The initial design of the PM systems was to collect data on every beam dump, e.g. every 8 hours under normal operation conditions, or multiple times an hour when problems arise. The newer use cases call for real-time constraints. The analysis result for XPOC must be delivered within 10 seconds, an SPS cycle analysis must be given within 7 seconds in order to use the result on the next cycle. A LINAC4 cycle lasts about 1 seconds so an analysis that interlocks the next cycle requires real-time constraints for the data transfer protocol and for the analysis framework.

In its current design, the PM system provides a good vertical scalability, meaning that resources can be added to the nodes with minimum downtime and impact on the service availability. This way it could satisfy the original use case and withstand the first extensions. The increase in data throughput to the PM system and the SPS use case integration both demonstrated that its horizontal scalability must be improved to sustain all its current and future use cases as well as to accommodate larger events and analysis processes.

The first Section presents the current functional design of the Post Mortem service. The second Section presents the architecture of the PM system to support the functional design and identifies the current shortcomings. The third Section presents the improvements we propose, the steps that are planned to achieve the migration, and the first results of the migration.

FUNCTIONAL BEHAVIOR

The LHC control devices are continuously recording information about their current state in a rolling buffer which they are ready to dump on demand from an external event, or on their own trigger. The data dump is identified by the device information and by an “event time stamp” given by the device. Typical events are beam dump, powering faults, self triggers, etc. If the data dump could not be performed completely and the data were not serialized to the permanent storage despite the fallback mechanism, the device is warned so that it can store the data locally on the front-end controller storage and try to dump them later on.

All the collected data dumps are forwarded to the Post Mortem event-builder. The event-building is data driven. When it receives a first data dump coming from a control device, it sets a data collection time window. All the subsequent data dumps arriving during this time window with an event time stamp close to the one of the first one will be associated to the event. The event-builder identifies and characterizes the type of the event according to the data dump patterns. If the event-builder only received data from power converters and quench protection systems, it is considered a powering event. If it received data from every LHC system, this is a global event which is in average 190 MB.

The event is then analyzed by a hierarchy of analysis modules. Modules can be dedicated to the analysis of a single type of system, or they can perform correlation between multiple systems.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

THE UNIFIED ANKA ARCHIVING SYSTEM – A POWERFUL WRAPPER TO SCADA SYSTEMS LIKE TANGO AND WINCC OA

D. Haas, S. Chilingaryan, A. Kopmann, D. Ressmann, W. Mexner, KIT, Karlsruhe, Germany

Abstract

ANKA realized a new unified archiving system for the typical synchrotron control systems by integrating their logging databases into the "Advanced Data Extraction Infrastructure" (ADEI). ANKA's control system environment is heterogeneous: some devices are integrated into the Tango archiving system, other sensors are logged by the Supervisory Control and Data Acquisition (SCADA) system WinCC OA. For both systems modules exist to configure the pool of sensors to be archived in the individual control system databases. ADEI has been developed to provide a unified data access layer for large time-series data sets. It supports internal data processing, caching, data aggregation and fast visualization in the web. Intelligent caching strategies ensure fast access even to huge data sets stored in the attached data sources like SQL databases. With its data abstraction layer the new ANKA archiving system is the foundation for automated monitoring while keeping the freedom to integrate nearly any control system flavor. The ANKA archiving system has been introduced successfully at three beamlines. It is operating stable since about one year and it is intended to extend it to the whole facility.

INTRODUCTION

ANKA is a third generation synchrotron light source operated by the Karlsruhe Institute of Technology (KIT). ANKA is operating seventeen beamlines and three more are under construction.

The control system of the ANKA beamlines is based on two independent components. The beamline instrumentation like the pressure, temperature or vacuum system is integrated in the SCADA system WinCC OA [1]. The experimental control, which includes movement of motors, triggering detectors et cetera is managed by Tango [2].

At ANKA the number of analogue sensors for long term archival per beamline can be estimated to 175 in WinCC OA plus 100 in Tango. Considering all 20 beamlines, these results to a number of approximately 5500 sensors, which have to be logged and archived at ANKA. Both systems have a separate solution to visualize the data for the user.

The main goal for the new unified ANKA archiving system is an easy and flexible configuration tool to select the data points for archival. And secondly to provide an unified fast graphical data browser [3]. Therefore this project splits in two parts. The first part was to interface the existing separate archiving solutions and the second part is to access all logged data of a beamline and the experiment. This data should be presented and retrieved in a modern, state-of-the-art web interface with fast data access. This offers the user a convenient way to analyse the data. ADEI [5], a web based interface for database query, developed by the Institute for Data Processing and Electronics (IPE) at KIT, fulfils exactly these requirements. Using ADEI as a viewer, respectively analysis tool and connecting it to the databases of the Tango archiving system and the WinCC OA Archiver plugin creates a platform to track and monitor the status of a beamline.

Due to the structure and available tools of this project, the main challenge was to combine and connect the three different parts, consisting of the Tango Archiving system, the WinCC OA History Database and ADEI.

INTEGRATION OF CONTROL SYSTEM HISTORY ARCHIVES

The integration of all separated archiving system in an overall distributed archiving system requires a TCP/IP based database interface of all sub systems, preferably SQL.

The Tango archiving system [4] allows logging of all Tango-attributes in a MySQL database. In 2014 this archiving system was already evaluated and implemented at ANKA to log the data of the experiment. The database of the SCADA part based on WinCC OA has no TCP/IP based SQL interface and is only accessible via a proprietary programming interface (API). Therefore ANKA developed a WinCC OA archiver plugin to record the logging data of preselected WinCC OA data points into a common SQL database. As the Tango Archiving System has been stored its logged data already into a MySQL database, we decided to use MySQL for the whole project. We found that the effort for the second option, to implement a new SQL interface to the existing C++ History API of WinCC OA, was too high.

The existing Tango Archiving System provides different advantages for logging the data points of the Tango control system. This system is a well-known and maintained archiving tool, which also includes with Mambo an easy to use graphical user interface.

Archiving the data-points of the SCADA system WinCC OA in a MySQL database required a new development. As the internal C like script language of WinCC OA is quite powerful, it was possible to use it to write all updated values directly to MySQL. For a convenient graphical user interface to configure archiving of the analogue values to be logged (Figure 2) a WinCC OA panel was developed. This GUI also allows different modes of logging. The data can be archived in different time periods or can be logged only if the value is changing.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

INTEGRATED DETECTOR CONTROL AND CALIBRATION PROCESSING AT THE EUROPEAN XFEL

A. Münnich, S. Hauf, B.C. Heisen, F. Januschek, M. Kuster, P.M. Lang, N. Raab,
T. Rüter, J. Sztuk-Dambietz, M. Turcato
European XFEL GmbH, Albert-Einstein-Ring 19, 22761 Hamburg, Germany

Abstract

The European X-ray Free Electron Laser is a high-intensity X-ray light source currently being constructed in the area of Hamburg, that will provide spatially coherent X-rays in the energy range between 0.25 keV and 25 keV. The machine will deliver 10 trains/s, consisting of up to 2700 pulses, with a 4.5 MHz repetition rate. The LPD, DSSC and AGIPD detectors are being developed to provide high dynamic-range Mpixel imaging capabilities at the mentioned repetition rates. A consequence of these detector characteristics is that they generate raw data volumes of up to 15 Gbyte/s. In addition the detector's on-sensor memory-cell and multi-/non-linear gain architectures pose unique challenges in data correction and calibration, requiring online access to operating conditions and control settings. We present how these challenges are addressed within XFEL's control and analysis framework Karabo, which integrates access to hardware conditions, acquisition settings (also using macros) and distributed computing. Implementation of control and calibration software is mainly in Python, using self-optimizing (py) CUDA code, numpy and iPython parallels to achieve near-real time performance for calibration application.

INTRODUCTION

The European X-ray Free Electron Laser (XFEL.EU) is an international research facility currently under construction in the area of Hamburg, Germany, which will start operation at the end of 2016 [1]. The superconducting linear accelerator of the facility will deliver electron bunches with an energy of up to 17.5 GeV, arranged in trains of typically 2700 bunches at a repetition rate of 4.5 MHz. Each train will be followed by a gap of 99.4 ms, during which data is read out from the detector front-end. The X-ray pulses will be particularly intense and ultra-short, down to a pulse length of less than 100 fs with a peak brilliance of 10^{33} photons/s/mm^2/mrad^2 per 0.1 % bandwidth [2].

The high repetition rate and the high peak brilliance pose unique challenges in terms of performance and radiation tolerance on the X-ray detector systems employed at the facility. As a result, each scientific instrument uses a detector optimized for its specific needs in terms of energy range, resolution, operation conditions, etc. Three large 2D area imaging detectors are being developed: the Adaptive Gain Integrating Pixel Detector AGIPD [3], DePFET Sensor with Signal Compression DSSC [4, 5] and a Large Pixel Detector LPD [6], which can run at the high repetition rate of 4.5 MHz. Smaller detectors like pnCCDs [7, 8] and Fast CCDs [9, 10] can be operated at 10 Hz. In addition, detector technologies such as Silicon Drift Detectors (SDDs), Micro-Channel Plates (MCPs) or silicon strip detectors [11] are being evaluated for specialized applications.

The large data volumes produced by the MHz-rate 2D detectors (10 to 15 GB per second per detector) and the huge number of calibration parameters (of the order of 10^9) require new concepts in data handling, calibration and processing in an integrated manner, aspects of which we will present.

THE KARABO DETECTOR CONTROL SOFTWARE FRAMEWORK

The software framework Karabo [12] will be used to control and manage the photon beamlines at the European XFEL as well as components within the experimental hutches including the detectors. Karabo's generic functionality covers all aspects from data acquisition (DAQ) over equipment control to data analysis. Figure 1 illustrates the different component categories available in the framework and their interaction via a JMS message broker or direct peer-to-peer (p2p) interfaces. In the following we present examples of the usage in the aforementioned categories: the control of a pnCCD-detector setup including the associated vacuum systems employing device composition, the control, DAQ and data-processing of an SDD-detector as well as pipeline-data analysis in the context of detector data correction and calibration.

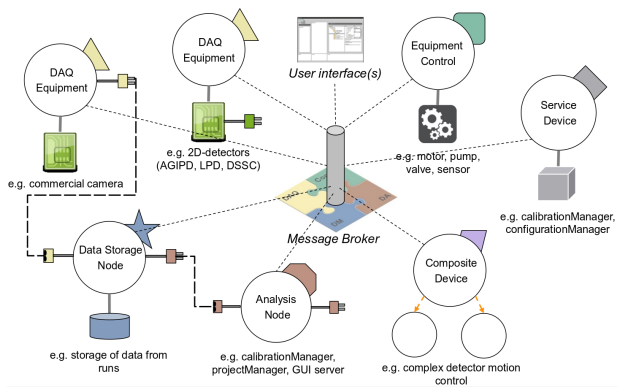


Figure 1: Schematic overview of Karabo's design concept. Functionality is implemented into the framework as devices (either `C++` or `Python3`), which can communicate either via a JMS message broker (shown here as monolithic, but in fact supporting dynamic fail-over and clustering) or p2p interfaces.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

DEVELOPMENT OF THE J-PARC TIME-SERIES DATA ARCHIVER USING A DISTRIBUTED DATABASE SYSTEM, II

N. Kikuzawa#, H. Ikeda, Y. Kato J-PARC, Tokai-mura, Naka-gun, Ibaraki, Japan
A. Yoshii, NS Solutions Corporation, Shinkawa, Chuo-ku, Tokyo, Japan

Abstract

J-PARC (Japan Proton Accelerator Research Complex) consists of Linac, 3 GeV rapid cycling synchrotron ring (RCS) and Main Ring (MR). In the Linac and the RCS, data of about 64,000 EPICS records have been acquired for control of these equipment. The data volume is about 2 TB every year, and the total data volume stored has reached over 12 TB. The data have been being stored by a Relational Database (RDB) system using PostgreSQL since 2006 in PostgreSQL, but it is becoming that PostgreSQL is not enough in availability, performance, and flexibility for our increasing data volume. In order to deal with these problems, we proposed a next-generation archive system using Apache Hadoop, a distributed processing framework, and Apache HBase, a distributed database.

In this paper we are reporting that we have re-designed and re-constructed the cluster with resolving some issues, including enhancing hardware of master nodes, creating scripts to automatically construct nodes, and introducing monitoring tools for nodes. Having adjusted the configurations of HBase/Hadoop and measured the performance of our new system, we are also reporting its results and considerations.

INTRODUCTION

J-PARC is controlled with a lot of equipment, and we have been archiving a time series of operation data provided from about 64,000 EPICS records for the Linac and the RCS since 2006 [1]. PostgreSQL has been used in the present data archiving system, but it has some problems of capacity, extensibility, and data migration. In order to deal with these problems, we proposed a next-generation archive system [2][3] using Apache Hadoop [4], a distributed processing framework, and Apache HBase [5], a distributed database.

In the previous paper we reported that we updated the versions of HBase/Hadoop composing our test system [6], to conquer a single point of failure by making redundant a master node, and we showed issues to fix our tools in the new system. Having adjusted the configurations of HBase/Hadoop and measured the performance of our new system, we are also reporting its results and considerations.

RECONSTRUCTION OF THE CLUSTER

Replacement Master Node

We have updated Hadoop to the version 2.2.0 now. Hadoop 2.x provides a hot standby NameNode, which can take over the state that the previous active NameNode has provided, with no downtime. At least three master nodes are needed for this function to deploy ZooKeeper [7] and JournalNode. ZooKeeper is a high-performance coordination service for distributed applications, and both Hadoop and HBase depend on. JournalNode is one of the components of Hadoop. ZooKeeper and JournalNode are based on a majority decision among nodes, and it is meaningful to deploy them on an odd number of machines. We replaced the master node 2 and 3. The layout of the system is illustrated in Figure 1 and the spec of our system is listed in Table 1.

When the master nodes has been replaced, processes of Hadoop has been relocated, and memory allocation has been reconsidered. The allocation of memory shows in Table 2.

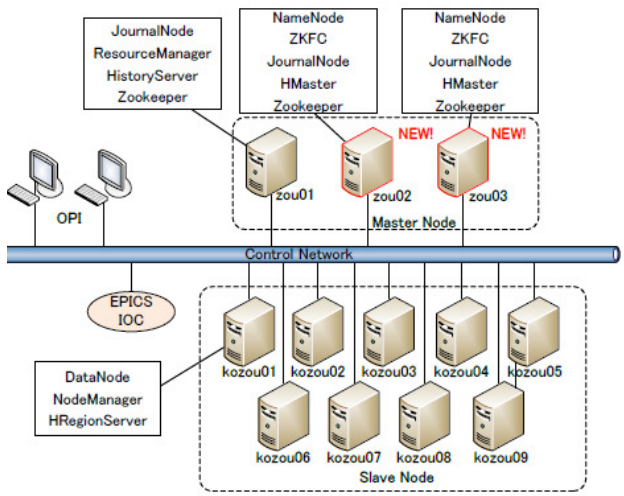


Figure 1: The archiving system configuration.

Table 1: Spec of Hadoop and HBase System

Master Node 1	DELL PowerEdge R610 CPU: Intel Xeon E5620 (4Core, 2.4GHz) MEM: 24GB HDD: 600GB x4 (RAID10)
Master Node 2, 3	DELL PowerEdge R320 CPU: Intel Xeon E5-1410v2 (4Core, 2.8GHz) MEM: 24GB HDD: 600GB x4 (RAID10)
Slave Node	DELL PowerEdge R410 CPU: Intel Xeon E5620 (4Core, 2.4GHz) MEM: 24GB HDD: 2TB x4
Software	OS: CentOS6.6 Hadoop: 2.2.0 HBase: 0.96..1.1 ZooKeeper: 3.4.5

#kikuzawa.nobuhiro@jaea.go.jp

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

MONITORING AND CATALOGUING THE PROGRESS OF SYNCHROTRON EXPERIMENTS, DATA REDUCTION, AND DATA ANALYSIS AT DIAMOND LIGHT SOURCE FROM A USER'S PERSPECTIVE

J. Aishima, A. W. Ashton, S. J. Fisher, K. E. Levik, G. Winter, Diamond Light Source Ltd, Didcot, Oxfordshire, UK

Abstract

The high data rates produced by the latest generation of detectors, more efficient sample handling hardware and ever more remote users of the beamlines at Diamond Light Source require advanced data reduction and data analysis tools to maximize the potential benefit to scientists. Here we describe some of our experiment data reduction and analysis steps, including real time image analysis with DIALS, our Fast DP and xia2-based data reduction pipelines, and Fast EP phasing and DIMPLE difference map calculation pipelines that aim to rapidly provide feedback about the recently completed measurements. SynchWeb, an interface to an open source laboratory information management system called ISPyB (co-developed at Diamond and the ESRF), provides a modern, flexible framework for managing samples and visualizing the data from all of these experiments and analyses, including plots, images, and tables of the analysed and reduced data, as well as showing experimental metadata, sample information.

INTRODUCTION

The availability of fast, continuous readout detectors (e.g. Pilatus3, Dectris[1]), highly automated sample handling systems such as BART robots[2], and intense X-ray beams at synchrotron sources have been crucial in increasing the rate of sample throughput and diffraction data collection at modern macromolecular crystallography (MX) beamlines. Non-automated evaluation of the many thousands of diffraction images that are generated per measurement in such experiments is no longer feasible, necessitating automated data reduction and analysis techniques combined with a laboratory information management system (LIMS) (SynchWeb [3]) that provides the user with the information necessary to evaluate the progress of their experiment while at the beamline.

SAMPLE EVALUATION

A typical macromolecular crystallography experiment on a beamline begins with visually locating a single crystal or using grid scanning[4] to locate samples or subsamples whose diffracting ability are assessed directly. For experiments where heavy atoms are incorporated for phasing purposes[5], a fluorescence scan may be performed to determine the optimal parameters for anomalous data collections.

STRATEGY

To determine the data required for an experiment, a screening data collection is typically performed with images at 0, 45, and 90 degrees. EDNA[6], a Python pipeline wrapping underlying software, is used on a multi-core computing cluster to calculate optimal data collections for several situations including native, anomalous, gentle (for radiation-sensitive samples), single anomalous diffraction (SAD), and considering measured flux. If applicable to the beamline, kappa alignment parameters (using XOAlign[7]) may also be calculated. Results of the strategy calculations are available in SynchWeb (Figure 1) and the desired EDNA strategy may also be retrieved when setting up a data collection using the Generic Data Acquisition software used at Diamond (GDA) [8].

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

FLYSCAN: A FAST AND MULTI-TECHNIQUE DATA ACQUISITION PLATFORM FOR THE SOLEIL BEAMLINES

N. Leclercq, J. Berthault, F. Langlois, S. Le, S. Poirier, Control and Data Acquisition Group, Synchrotron SOLEIL, France
J. Bisou, F. Blache, Electronics Group, Synchrotron SOLEIL, France
K. Medjoubi, C. Mocuta, Scientific Division, Synchrotron SOLEIL, France

Abstract

Synchrotron SOLEIL [1] is continuously optimizing its 29 beamlines in order to propose state of the art synchrotron radiation based experimental techniques to its users. Among the topics addressed by the related transversal projects, the enhancement of the computing tools is identified as a high priority task. In this area, the aim is to optimize the beam time usage providing the users with a fast, simultaneous and multi-technique scanning platform. The concrete implementation of this general concept allows the users not only to acquire more data in the same amount of beam time, but also to address specific phenomena otherwise difficult to tackle (like processes involving rapid transformation of the sample).

This paper gives the reader a detailed overview of the so called 'Flyscan' platform. It notably details the solution retained to generalize a prototype previously developed as a proof of concept [2][3]. Some application examples are also reported.

MOTIVATION AND SPECIFICATIONS

The main motivation for launching the Flyscan project was beam time optimization. Acquiring data faster obviously means obtaining more exploitable data at constant beam time. But it also means that one will be able to quickly discover that experimental conditions will not provide usable data – a benefit which can potentially save hours of expensive and highly sought out beam time.

From a scientific point of view, fast data acquisition also gives access to dynamics of processes involving rapid transformation of the sample (e.g. thermal annealing, continuous mechanical deformation, etc.).

The technical translation of the requirement above led to the specification of an N-dimensional scanning platform whose first dimension is "continuous". Here continuous is opposed to the step-by-step approach – *i.e.* the classical scanning technique known to generate a huge amount of dead time (acceleration / deceleration times of the motors, steady state return information, etc.). In order to reinforce the concept of beam time optimization, the ambition is also to combine the continuous acquisition approach with a multi-technique solution on the sensors (*i.e.* detectors) side. Clearly speaking, the idea is to simultaneously collect – spatially or temporally – correlated data from different sensors along the continuous trajectory of one (or more) actuator(s).

DESIGN OVERVIEW

From a technical point of view, the Flyscan backbone relies on 3 buses: the timing bus, the communication bus and the data bus.

The timing bus is in charge of the synchronization signals (*i.e.* trigger) distribution. Timing events are polymorphic and can be dispatched in the form of hardware (*e.g.* TTL) or software (*e.g.* UDP) notifications. The communication bus routes the control oriented exchanges between the scan actors. Since the latter are implemented as Tango devices [4], these exchanges are mainly composed of Tango queries and events emitted over an Ethernet network. Finally, the data bus transports the experimental data from their respective source to a central point where they are merged into a common repository. The choice has been made to use HDF5 [5] files for both data transport and storage. Figure 1 illustrates the Flyscan architecture.

It results that the Flyscan system can be seen as an aggregation of individual components sharing a common timebase to guarantee temporally – and consequently, spatial - correlation of data produced by heterogeneous sensors distributed over a network.

IMPLEMENTATION DETAILS

Timing

In the Flyscan scheme, everything starts with a software or hardware master clock. The master clock source selection is guided by application requirements. For continuous motion driven experiments – such as tomography or X-ray microscopy – a hardware solution, named *Spietbox*, has been developed to generate a spatially periodic pattern (*e.g.* one clock pulse every tenth of a degree of a rotation) [6]. The *Spietbox* can be coupled with a counting board in order to sample and record the exact axis positions at which data is acquired on each data source while the axis is moving. Solutions also exist to generate slave clocks.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

THE AUSTRALIAN STORE.SYNCHROTRON DATA MANAGEMENT SERVICE FOR MACROMOLECULAR CRYSTALLOGRAPHY

Grischa R Meyer, Steve Androulakis, Philip Bertling, Ashley Buckle,
Wojtek James Goscinski, David Groenewegen, Chris Hines, Anitha Kannan,
Sheena McGowan, Stevan Quenette, Jason Rigby, Patrick Splawa-Neyman,
James M Wettenhall, Monash University, Clayton, Australia
David Aragao, Tom Caradoc-Davies, Nathan Mudie, Synchrotron Light Source Australia, Clayton
Charles Bond, University of Western Australia, Crawley, Australia

Abstract

Store.Synchrotron is a service for management and publication of diffraction data from the macromolecular crystallography (MX) beamlines of the Australian Synchrotron. Since the start of the development, in 2013, the service has handled over 51.8 TB of raw data (~ 4.1 million files). Raw data and autoprocessing results are made available securely via the web and SFTP so experimenters can sync it to their labs for further analysis. With the goal of becoming a large public repository of raw diffraction data, a guided publishing workflow which optionally captures discipline specific information was built. The MX-specific workflow links PDB coordinates from the PDB to raw data. An optionally embargoed DOI is created for convenient citation. This repository will be a valuable tool for crystallography software developers. To support complex projects, integration of other instruments such as microscopes is underway. We developed an application that captures any data from instrument computers, enabling centralised data management without the need for custom ingestion workflows. The next step is to integrate the hosted data with interactive processing and analysis tools on virtual desktops.

INTRODUCTION

The deposition and the wide availability of raw diffraction data have far-reaching benefits for the structural biology community [1–4]. We recently created TARDIS, a suite of tools for the deposition of X-ray diffraction images in an open access repository to facilitate their deposition using federated institutional repositories [3]. Subsequent engagement with IUCR and closer relationships with the Australian Synchrotron prompted us to develop a new framework for raw X-ray data archival and dissemination — the *Store.Synchrotron* service, which is described in this manuscript.

An earlier but more detailed description of this service is available at [5].

THE OPERATION OF STORE.SYNCHROTRON FOR THE AUSTRALIAN SYNCHROTRON

The *Store.Synchrotron* service (`https://store.synchrotron.org.au`) has been deployed to receive diffraction data from the macromolecular crystallography (MX) beamlines at the Australian Synchrotron (AS) automatically.

When a user starts collecting data at the beamline, all raw diffraction data and automated data processing results are transferred in real-time and on the next day respectively to *Store.Synchrotron* without user intervention.

This system is accessible through any web browser. Data is available immediately at time of collection to all authorised users via the internet. If desired, the data can be opened to the public.

Figure 1 shows a schematic overview of the service.

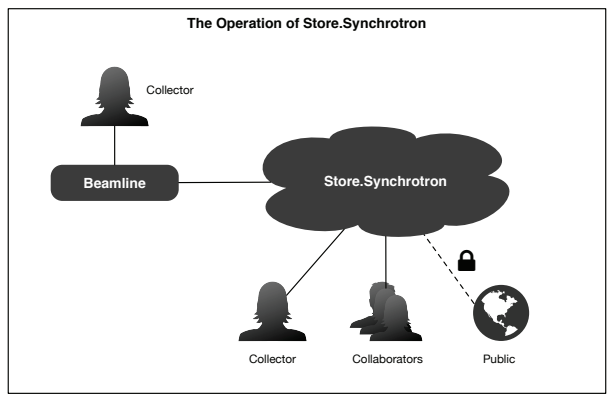


Figure 1: Schematic overview of *Store.Synchrotron*

The Need for an Automated Data Management System

At the Australian Synchrotron (AS) up until recently the most practicable solution for data management has been to take portable hard drives to the beamlines. The use of portable drives is a problem, however, because these drives may remain the lab's sole resource for raw data. After researchers have left, it may be difficult to correlate raw frames with datasets and structures, and a single disk is very prone to failure and data loss.

With *Store.Synchrotron* users now have for the first time a long-term, secure archive of their raw data, including metadata, which can be used to relate raw frames to projects and visits to the Australian Synchrotron.

In addition to the practical need for a robust data archive, new requirements by funding agencies for researchers to have a data management plan are conveniently solved by a

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

A DATA MANAGEMENT INFRASTRUCTURE FOR NEUTRON SCATTERING EXPERIMENTS IN J-PARC/MLF

K.Moriyama# and T.Nakatani, J-PARC, Ibaraki, Japan

Abstract

The role of data management is one of the greatest contributions in the research workflow for scientific experiments such as neutron scattering. The facility is required to safely and efficiently manage a huge amount of data over the long duration, and provide an effective data access for facility users promoting the creation of scientific results. In order to meet these requirements, we are operating and updating a data management infrastructure in J-PAPC/MLF, which consists of the web-based integrated data management system called the MLF Experimental Database (MLF EXP-DB), the hierarchical raw data repository composed of distributed storages, and the integrated authentication system. The MLF EXP-DB creates experimental data catalogues in which raw data, measurement logs, and other contextual information on sample, experimental proposal, investigator, etc. are interrelated. This system conducts the reposition, archive and on-demand retrieve of raw data in the repository. Facility users are able to access the experimental data via a web portal. This contribution presents the overview of our data management infrastructure, and the recent updated features for high availability, scaling-out, and flexible data retrieval in the MLF EXP-DB.

INTRODUCTION

The Materials and Life Science Experimental Facility (MLF) at the Japan Proton Accelerator Research Complex (J-PARC) is the experimental facility for neutron scattering, providing domestic and international users in a wide variety of research fields with one of the highest intensity pulsed neutron beam in the world. Twenty neutron experimental instruments equipped with large-area neutron detectors and various sample environmental apparatuses are currently in operation. MLF annually supplies beamtime of two hundreds days for about a thousand users and about several hundreds of experimental proposals are performed using these neutron instruments.

Neutron instruments create a huge amount of experimental data. A typical volume of data produced in one experiment is from several tens MB to several hundreds GB. Since a lot of experiments under various experimental conditions are carried out in a short period of time using high intensity beam, the total amount of data generated annually in a whole facility at full performance is on the order of PB.

In those situations, it is required to safely and efficiently manage a huge amount of data over the long duration. In particular, an appropriate protection and preservation of the experimental data as intellectual property produced at a large experimental facility such as MLF is required recently. Furthermore, it is required to provide facility users with a data access service enabling effective utilization of a lot of experimental data. To meet these requirements, we are operating an infrastructure for a central data management and access in the whole facility. In addition, we are updating the MLF EXP-DB, which is the core system of our infrastructure, responding to the expansion and growth of experimental environment based on an improvement plan in the facility. Since the neutron beam currently provided with bombarding of 500kW proton beam is finally increasing up to the intensity equivalent to the power of 1 MW proton beam, therefore the infrastructure should respond the enhancement of data production rate due to the intensity enhancement of neutron beam.

#kentaro.moriyama@j-parc.jp

DATA MANAGEMENT INFRASTRUCTURE

A data management infrastructure in J-PARC/MLF is comprised of several database systems and the data repository. The schematic overview of the infrastructure is shown in Fig. 1. The MLF EXP-DB is the core system of the infrastructure, and responsible for managing data and providing data access. This system catalogues experimental data and manages measurement raw data centrally in the whole facility. Measurement raw data is preserved in the data repository. The MLF Business Database (MLF BIZ-DB) collects application information such as experimental proposal, primary investigator, and sample from the front end system of facility users. Also, there are an integrated authentication system (User DB), which manages user credentials of systems comprising infrastructure, and a sample database (Sample-DB) for a safety management of handling samples. The MLF EXP-DB collects automatically every variety of information related to experiment by interacting with these database systems.

In local system of neutron instrument at MLF, the MLF control software framework "IROHA" [1], which is adopted as standard software for measurement, coordinates measurement with controlling data acquisition and sample environmental devices. In response to the message from IROHA, the data acquisition system (DAQ) produces raw data files into the data storage. Also, measurement log file including measurement conditions such as sample device configuration and results is created as "experimental meta-data" in XML-format by IROHA. This experimental

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

BEAM TRAIL TRACKING AT FERMILAB

Dennis J. Nicklaus, Linden Ralph Carmichael, Richard Neswold, Zongwei Yuan.
Fermilab, Batavia, IL 60510, USA

Abstract

We present a system for acquiring and sorting data from select devices depending on the destination of each particular beam pulse in the Fermilab accelerator chain. The 15 Hz beam that begins in the Fermilab ion source can be directed to a variety of additional accelerators, beam lines, beam dumps, and experiments. We have implemented a data acquisition system that senses the destination of each pulse and reads the appropriate beam intensity devices so that profiles of the beam can be stored and analysed for each type of beam trail. We envision utilizing this data long term to identify trends in the performance of the accelerators.

INTRODUCTION

The Fermilab particle beam starts in its ion source, and from there it can be directed to a wide variety of beamlines and experiments. This paper describes the software created to track the accelerator efficiencies (amount of beam transported without loss) along the various beamlines and for the designated destinations. We use the term trail to designate each of the potential uses and destinations.

Trails can have a wide range of complexity. The simplest trails consist of only one 15Hz Linac beam pulse, and might go to a local beam dump, or be used to study Linac performance. The more complicated pulses involve multiple 15Hz pulses from the Linac, and involve the Linac, Booster, Recycler, and Main Injector accelerators as shown in Fig. 1. Final destinations include various beam dumps along the way, a neutron therapy target, short and long baseline neutrino beamlines, the muon campus experiments, or a test beam area.

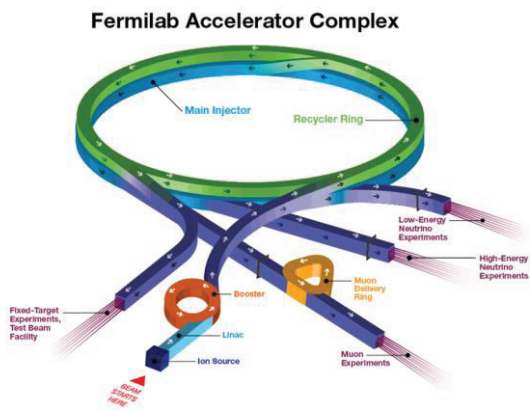


Figure 1: The Fermilab Accelerator Complex.

Collection of these more complex trails is also complicated by the fact that some of the simpler 15Hz pulse trails may occur while the more complicated trail is still underway in the downstream accelerators.

The process starts with the overall beam pulse planning controlled by Fermilab's Sequencer, which coordinates the production of timelines generated by the Time Line Generator (TLG) with other accelerator activities. At the implementation level, a timeline is a series of private bus events (TCLK). Various accelerator components are programmed to respond to these events to send the beam to the desired location. These systems have long been in routine use for accelerator operation.

For this beam trail tracking, we added or enhanced several software programs. To start the process, the TLG now encodes the trail number onto our MDAT data bus when it begins instructions for each beam pulse. The Beam Cycle Coordinator (BCC) decodes that trail number and makes it available to our Acnet control system, in addition to many digital status bits that indicate readiness of various beamlines, components, permits, and interlocks. These status indicators also show whether the accelerator chain was ready and beam pulse had the opportunity to follow the indicated trail, or whether it had to be aborted.

With the above groundwork laid, the main implementation for beam trail tracking begins. The tracking software has three main components:

- The Correlator, which collects data from Acnet-connected devices, groups together readings from the same pulse, and sorts them by trail number.
- The Acnet Formatter, which takes the data from the Correlator, aggregates and re-formats it and makes it available to Acnet.
- The Database Interface, which reads the data over Acnet and stores it into a relational database. Along the way, this also calculates various sums and means, and we have developed a user interface to help retrieve data of interest.

BACKGROUND

Main Injector Ramp Cycle

In order to understand some of the requirements of the trail tracking software, it is useful to have a basic understanding of our Main Injector ramp cycle. The Main Injector accelerates protons from 8 GeV to 120 GeV, taking a little over one second for this acceleration. While it is ramping up and down, multiple batches of protons from the Booster can be injected into the Recycler

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

PROCESSING HIGH-BANDWIDTH BUNCH-BY-BUNCH OBSERVATION DATA FROM THE RF AND TRANSVERSE DAMPER SYSTEMS OF THE LHC

M. Ojeda, P. Baudrenghien, A.C. Butterworth, J. Galindo, W. Hofle,
T.E. Levens, J.C. Molendijk, D. Valuch, CERN, Geneva, Switzerland
F. Vaga, University of Pavia, Pavia, Italy

Abstract

The radiofrequency and transverse damper feedback systems of the Large Hadron Collider digitize beam phase and position measurements at the bunch repetition rate of 40 MHz. Embedded memory buffers allow a few milliseconds of full rate bunch-by-bunch data to be retrieved over the VME bus for diagnostic purposes, but experience during LHC Run I has shown that for beam studies much longer data records are desirable. A new "observation box" diagnostic system is being developed which parasitically captures data streamed directly out of the feedback hardware into a Linux server through an optical fiber link, and permits processing and buffering of full rate data for around one minute. The system will be connected to an LHC-wide trigger network for detection of beam instabilities, which allows efficient capture of signals from the onset of beam instability events. The data will be made available for analysis by client applications through interfaces which are exposed as standard equipment devices within CERN's controls framework. It is also foreseen to perform online Fourier analysis of transverse position data inside the observation box using GPUs with the aim of extracting betatron tune signals.

SYSTEM OVERVIEW

The "Observation Box" (ObsBox) is a custom acquisition system designed to overcome the bandwidth limits of the VME bus to capture diagnostic data. Figure 1 shows an overview of the system.

The transmitters shown in the figure are typically custom-built VME modules in the beam feedback systems with embedded FPGAs that sample different aspects of the accelerator's beam at a high frequency. In many cases, the acquisition of these data streams for diagnostic purposes can be done through the VME bus into the Linux host residing in the crate. However, some analyses of the beam data can really benefit from a substantially higher data bandwidth than the VME bus can provide.

Typically, the transmitters are close to the beam, in the underground caverns. For accessibility and maintenance reasons, the receiver, a Linux server, is installed on the surface, with a fiber link running between the two.

The new transmitters designed for the ObsBox system stream the data out through the fiber links. At the server end, for each fiber, a CERN-designed PCI Express v1 module is used to act as the receiver: the Simple PCI Express Carrier (SPEC) module [1].

The ObsBox server runs a customized Linux kernel intended for hosting real-time software. A tailored Linux kernel device driver is used to control the SPEC modules over the PCIe bus and is able to stream the data to a user-space process through syscalls and a *sysfs* interface.

The user-space process is responsible for reading the driver streamed data regularly and quickly enough and storing it into very large memory buffers, where the data waits to be requested eventually by clients. Clients connect to the ObsBox server over the CERN's technical network, a standard Ethernet network.

The user-space process is also foreseen to perform online data analysis if the network bandwidth is not enough to stream it out to clients. Since at this stage the data is local, using a coprocessor (like a GPU) to perform the analysis is a viable option, which in turn opens the possibility of running other kinds of analysis for which the CPU may not be performant enough.

TRANSMITTER (VME MODULE)

The Longitudinal Position Measurement VME module (LPM) comprises an analogue front-end and a digital signal processing FPGA platform generating digitized information that can be transmitted via fast digital serial links for further processing. The hardware is based on the beam phase measurement module originally developed for the LHC RF Beam Control [2].

Two signals are acquired in the RF front-end: the beam pickup and the 400.8 MHz RF voltage vector sum of eight cavities. These are demodulated into analog I/Q signals and bunch synchronously sampled by four 14-bit AD converters to give I/Q pairs at a data rate of 40.08 MHz. In the FPGA platform, the VME card is equipped with four gigabit links. The bunch synchronous data of the two I/Q signals (at a rate of 40.08 MHz) is "repackaged" into 50 MHz serial data frames, computing 1 Gbps per link. A FIFO synchronization approach has been used to handle the data crossing among clock domains.

The four data streams are sent over a copper to optical converter card towards the ObsBox. At the receiver side, the four optical links are resynchronized again with 40.08 MHz and the revolution frequency clock. Signal generators are also available within the FPGA platform enabling digital calibration and link debugging.

Similiarly, the LHC transverse damper system (ADT) uses two or four pickups per plane per beam to measure beam position at 40.08 Msamples/s, digitized in a VME beam

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

DEVELOPING HDF5 FOR THE SYNCHROTRON COMMUNITY

N. Rees, Diamond Light Source, Oxfordshire, UK
H. Billich, PSI, Villigen, Switzerland
A. Götz, ESRF, Grenoble, France
Q. Koziol & E. Pourmal, The HDF Group, Champaign, IL, USA
M. Rissi, Dectris AG, Baden, Switzerland
E. Wintersberger DESY, Hamburg, Germany

Abstract

HDF5[1] and NeXus[2] (which normally uses HDF5 as its underlying format) have been widely touted as a standard for storing Photon and Neutron data. They offer many advantages to other common formats and are widely used at many facilities. However, it has been found that the existing implementations of these standards have limited the performance of some recent detector systems. This paper describes how the synchrotron light source community has worked closely with The HDF Group to drive changes to the HDF5 software to make it more suitable for their environment. This includes developments managed by a detector manufacturer (Dectris - for direct chunk writes) as well as synchrotrons (DESY, ESRF and Diamond - for pluggable filters, Single Writer/Multiple Reader and Virtual Data Sets).

INTRODUCTION

The original initiative for this work came from a workshop held at Paul Scherrer Institute (PSI) in late May 2012. This was jointly organised by PSI and Dectris and brought together representatives from the synchrotron community and employees of The HDF Group. A number of issues were identified at that workshop, including:

1. The single threaded nature of the HDF5 library made it difficult to benefit from the multicore architecture of modern CPU's.
2. Compression was an important mechanism to improve data throughput in many synchrotron use cases, but including any form of compression limited the throughput of the library and also could not be used when writing in parallel with the Parallel HDF5 library (pHDF5).
3. Use of a hierarchical format like HDF5 led to the need to store many detector frames in one file. However, this increases the latency of any processing because HDF5 files could not be open for read and write simultaneously.

Since this meeting the authors have collaborated to fund a number HDF5 developments that have eliminated many of these problems for practical synchrotron use cases.

DIRECT CHUNK WRITES

This development was sponsored by Dectris and PSI and introduced a new function `H5DOwrite_chunk` in the high-level HDF5 library which bypasses the data gathering and scattering, data conversion, filter pipeline, and chunk cache and writes the data chunk to the file directly (see Fig. 1). This allows the user program to compress the data outside the library – potentially using parallel algorithms or hardware accelerators. It also avoids a number of data copies, which limits any dataflow through the filter pipeline to ~500 MB/sec on typical processors.

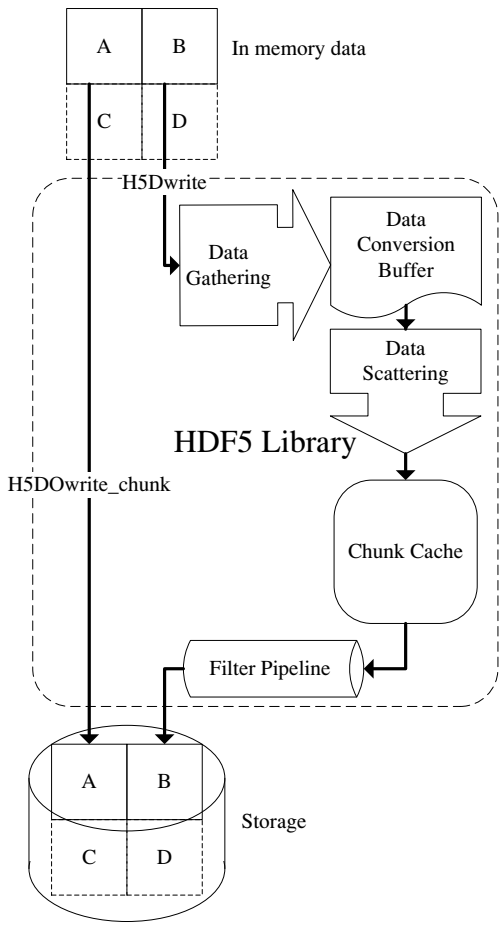


Figure 1: HDF5 data writing, showing the elements that `H5DOwrite_chunk` bypasses.

The feature was released in HDF5 version 1.8.11 in May 2013. The implementation and testing required about 5 months of work and, while the function did not present any technical difficulties, fitting it into the HDF5 library required some rethinking of the HDF5 source code architecture. The latest versions of the library are set for

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

ILLUSTRATE THE FLOW OF MONITORING DATA THROUGH THE MeerKAT TELESCOPE CONTROL SOFTWARE

M. Slabber*, SKA SA, Cape Town, South Africa
M.T. Ockards†, SKA SA, Cape Town, South Africa

Abstract

The MeerKAT telescope [1], under construction in South Africa, is comprised of a large set of elements. The elements expose various sensors to the Control and Monitoring (CAM) system, and the sampling strategy set by CAM per sensor varies from several samples a second to infrequent updates. This creates a substantial volume of sensor data that needs to be stored and made available for analysis. We depict the flow of sensor data through the CAM system, showing the various memory buffers, temporary disk storage and mechanisms to permanently store the data in HDF5 format on the network attached storage (NAS).

ELEMENTS THAT MAKE UP CAM

Sensor

Sensors are either physical sensors on devices, aggregated/calculated or internal metrics from software components. Some examples of physical sensors are tilt, temperature, wind speed and motor rpm. It is estimated that once MeerKAT is completed there will be close to 80 000 sensors in the system.

Device

Devices are the hardware that constitute the telescope. Devices are engineered to have Ethernet interfaces and can communicate via Karoo Array Telescope Communication Protocol (KATCP). Where devices do not have KATCP interfaces, translators are created [2]. The devices expose sensors and requests on the KATCP interface [3].

Component

CAM consists of many distinct components. Components are the standalone building blocks of the CAM system. Interaction with components are done with the KATCP protocol. The components expose sensors and requests over KATCP.

Node

Nodes are virtual containers on the host hypervisor [4]. The MeerKAT nodes run Ubuntu 14.04 LTS Linux operating system. When MeerKAT is fully deployed, it is expected to have 13 nodes.

Proxy

Proxy components make sensors and requests on devices available to the CAM system. Sensors and requests from the devices can be discarded or even rolled up into new sensors and requests on the proxy component. The proxy component also exposes sensors and requests of its own.

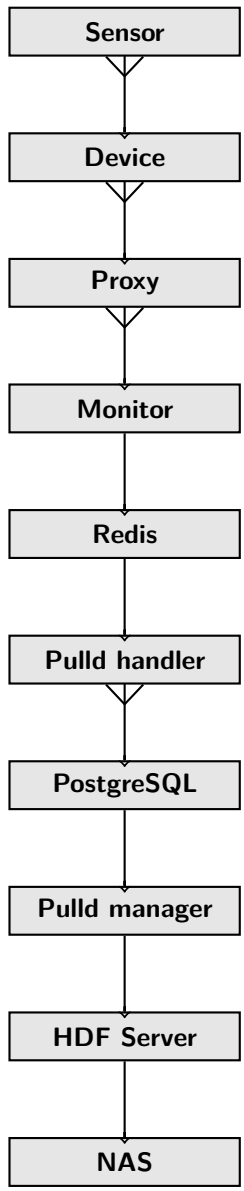


Figure 1: The progression of storing a sample, from sensor to archive.

Controller

Controller components connect to other subsystems. As with a proxy component, sensors and requests from the subsystem can be exposed or rolled up into new sensors and requests. Controllers always use KATCP to communicate to

* martin@ska.ac.za
† tockards@ska.ac.za

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

A SYSTEMATIC MEASUREMENT ANALYZER FOR LHC OPERATIONAL DATA

G. Valentino*, X. Buffat, D. Kirchner, S. Redaelli, CERN, Geneva, Switzerland

Abstract

The CERN Accelerator Logging Service stores data from hundreds of thousands of parameters and measurements, mostly from the Large Hadron Collider (LHC). The systematic measurement analyzer is a Java-based tool that is used to visualize and analyze various beam measurement data over multiple fills and time intervals during the operational cycle, such as ramp or squeeze. Statistical analysis and various manipulations of data are possible, including correlation with several machine parameters such as β^* and energy. Examples of analyses performed include checks of collimator positions, beam losses throughout the cycle and tune stability during the squeeze which is then used for feed-forward purposes.

INTRODUCTION

The CERN Large Hadron Collider (LHC) was designed and built to accelerate two counter-rotating beams to an energy of 7 TeV and deliver high energy collisions of protons or heavy ions in the four experiments, namely ATLAS, ALICE, CMS and LHCb [1]. Data from thousands of devices in the CERN accelerator infrastructure, amounting to around 1 million signals and 50 TB/year are stored by the CERN accelerator logging service (CALS) [2].

An overview of the software architecture is shown in Fig. 1. Two Oracle databases are used for storage. The Measurement database (MDB) is used to store raw data for a short period of time, which is then filtered and transferred to the Logging database (LDB) for indefinite storage of filtered data. The databases subscribe to specific properties of each device via the Java API for Parameter Control (JAPC) [3] or PVSS for the industrial SCADA systems.

A Java GUI called TIMBER is provided to visualize and extract the logged data. However, in order to assess the stability of the machine and perform analyses over several fills and machine modes, a tool called Systematic Measurement Analyzer was developed. This paper describes the implementation of the tool as well as some typical use case scenarios where it has proved to be useful.

LHC MACHINE CYCLE

The operation of the LHC follows well-established stages, which together form a machine cycle. A typical LHC machine cycle is shown in Fig. 2. At the injection stage (1), the LHC receives two beams from the Super Proton Synchrotron (SPS) at an energy of 450 GeV per beam. The beams arrive in bunch trains, which may consist of 1-144 bunches, depending on whether the machine is to be filled for beam tests or physics. When the filling procedure is completed, both beams are ramped up (2) until the desired energy is reached. At flat top (3), the machine operators initiate the beam squeeze procedure (4), in which corrector magnets are used to shrink the beam size at the experimental insertion points (IPs) to achieve the desired β^* (the β-function at the experimental IPs).

Up to this point, the beams are separated by several σ (r.m.s. beam size) in all IPs. Hence, the final step is to collapse the separation orbit bumps and bring the beams into collisions. The operational state known as stable beams (5) is declared, and the experiments begin taking data. The beams may be extracted or dumped (6) at any point during the fill due to operational requirements, equipment failures, beam instabilities or operator mistakes. When this happens, the machine is ramped down (7) to injection energy in preparation for the next fill. The data presented in Fig. 2 are taken from fill number 3131, in which a total of 1374 bunches were present in the machine, and $130 \times 130 \mu b^{-1}$ of luminosity was delivered to the ATLAS and CMS experiments in 3.5 hours of stable beams.

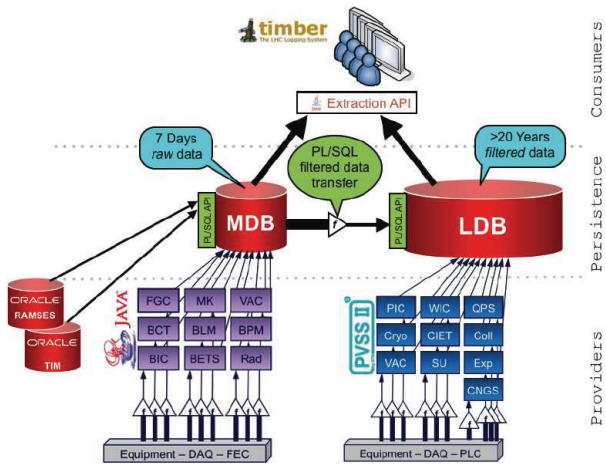


Figure 1: The CERN accelerator logging service architecture [2].

SOFTWARE ARCHITECTURE AND UI

The systematic measurement analyzer is written in Java, which is the standard for LHC operational applications run-

* gianluca.valentino@cern.ch

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

FORMALIZING EXPERT KNOWLEDGE IN ORDER TO ANALYSE CERN'S CONTROL SYSTEMS

A. Voitier, M. Gonzalez-Berges, F.M. Tilaro, CERN, Geneva, Switzerland
M. Roshchin, Siemens AG, Corporate Technology, München, Germany

Abstract

The automation infrastructure needed to reliably run CERN's accelerator complex and its experiments produces large and diverse amounts of data, besides physics data. Over 600 industrial control systems with about 45 million parameters store more than 100 terabytes of data per year. At the same time a large technical expertise in this domain is collected and formalized. The study is based on a set of use cases classified into three data analytics domains applicable to CERN's control systems: online monitoring, fault diagnosis and engineering support. A known root cause analysis concerning gas system alarms flooding was reproduced with Siemens' Smart Data technologies and its results were compared with a previous analysis. The new solution has been put in place as a tool supporting operators during breakdowns in a live production system. The effectiveness of this deployment suggests that these technologies can be applied to more cases. The intended goals would be to increase CERN's systems reliability and reduce analysis efforts from weeks to hours. It also ensures a more consistent approach for these analyses by harvesting a central expert knowledge base available at all times.

INTRODUCTION

CERN employs about 600 industrial Supervisory Control and Data Acquisition (SCADA) systems for the supervision and monitoring of its accelerators, detectors and infrastructure machines. While the day-to-day operations are running smoothly, a growing need appeared to exploit the data generated by these applications. Indeed, collectively they produce more than 100 terabytes of control data over 45 million parameters. The gathered data could be seen as a deep reflection of the current state of the processes under control. A lot of information about the performance, stability and overall behaviour of the machines resides within these data. Today, an expert can manually follow the signals deemed important and apply his/her knowledge to maintain a good level of service. However, current tools for industrial control systems are not properly designed for doing such dedicated analysis. The external analysis tools used today are not well integrated with the operator applications. Moreover the size and complexity of some systems does not allow running advanced data analytics methods in normal office computers. In addition, this way of working is a non-sense for analyses that have to be part of the operational tools themselves. Lastly, while the experts are knowledgeable about their domain, they are not properly skilled to perform these data analysis or computing problems tasks.

In this context, formalizing expert knowledge means capturing the methods and knowledge used by experts and transforming them into analyses that can be scaled out to all similar systems without burdening the users with a vast amount of manual operations to carry out.

From a computer science point of view, this problem is part of the Data Analytics, or Big Data, field which combines technologies for processing vast amount of data carrying out analytical tasks tailored in our case to industrial control system needs

Once an analysis can be applied to solve a control problem, it becomes part of the control system itself making possible to raise awareness of operators on specific issues from the supervision application they are accustomed to.

This paper will present the knowledge capture and the methods used to detect specific conditions traditionally done "by hand" (exporting data, importing them into a spreadsheet software or writing an adhoc script) or by looking at trends. Then, we discuss briefly their implementation as analytics tools that were then scaled out and applied automatically with the help of readily available software solutions.

CONTROL SYSTEM DATA ANALYSIS

Control systems play an essential role to run the CERN accelerators complex and generate a huge amount of data that can be used to analyse the behaviour of these systems and to find insights useful to operators and experts. The data analytics activities have been divided into three different categories:

- Online monitoring
- Fault diagnosis
- Engineering design

These three families of analysis focus on control data to offer analytical services as added value on top of the traditional industrial services.

Online Monitoring

Currently the monitoring and operation of CERN industrial systems is – for most parts – achieved through the deployment of specific applications based on the commercial SCADA software WinCC Open Architecture [1].

Therefore, these data analytics activities aim to add new features and services to this monitoring layer. This close integration allows reaching a high level of automation

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

INTEGRATING WEB-BASED USER INTERFACE WITHIN CERN'S INDUSTRIAL CONTROL SYSTEM INFRASTRUCTURE

A. Voitier, M. Gonzalez-Berges, P. Golonka, CERN, Geneva, Switzerland

Abstract

For decades the user interfaces of industrial control systems have been primarily based on native clients. However, the current IT trend is to have everything on the web. This can indeed bring some advantages such as easy deployment of applications, extending HMIs with turnkey web technologies, and apply to supervision interfaces the interaction model used on the web. However, this also brings its share of challenges: security management, ability to spread the load and scale out to many web clients, etc... In this paper, the architecture of the system that was devised at CERN to decouple the production WINCC-OA based supervision systems from the web frontend and the associated security implications are presented together with the transition strategy from legacy panels to full web pages using a stepwise replacement of widgets (e.g. visualization widgets) by their JavaScript counterpart. This evolution results in the on-going deployment of web-based supervision interfaces proposed to the operators as an alternative for comparison purposes.

INTRODUCTION

Supervisions for industrial control systems are applications that have to remains stable and shielded from any potential factor of disruption. As such, they are usually deployed on dedicated and isolated servers to minimize risks and control its access.

However, users of these systems have a pressing demand of accessing, or at list viewing, the state of their systems without the protective hurdles. Another limitation is the need to install specific software on the computer used for accessing remotely the supervision applications.

These requirements are following the tendency of the present computing era where we got accustomed to seamlessly access any remote services through a simple web-browser or a mobile application. A so-called Web UI (Web User Interface) for the SCADA tool used at CERN is implemented to answer these requests, provided security is strictly enforced.

This UI allows to view existing panels and synoptic used in the current native UI. The graphical engineering editor of the SCADA tool is also accessible with a web browser, suppressing the need to install any specific software. Another important advantage is to profit from standard and modern web technologies to enhance the usability of user interfaces as a wide range of versatile graphical widgets are available to web-enabled technologies [1].

For the end users it brings new paradigms of using supervision applications. It allows accessing them from a field intervention, or from a non-professional place for rapid remote diagnostics. Yet another usage is to broadcast the state of a particular machine to hundreds or thousands of users of this machine, which is a need for physics experiment collaborations.

This paper will present how this project went through several phases of iterative researches. First by prototyping a Web UI with classic web components. Then, facing important limitations, introducing an alternative approach similar to a remote desktop software embedded in a browser.

WEB ACCESS FOR INDUSTRIAL CONTROL SYSTEMS

At CERN, supervisions for industrial control applications amount to a total of 600 systems. All production applications are hosted on a dedicated network isolated from the intranet and internet by firewalls. To enable a web access, the following four main requirements have been postulated.

Security: Making control systems of very costly machines accessible to the entire world means security is paramount. A powerful access control has to be used while keeping the balance with usability. Indeed, with a very secure but hard to use system users have a tendency to circumvent the secure measures, putting the protected systems at risk. Another point is to stick to web security standards extensively reviewed and tested by experts and not forge custom solutions.

Legacy: In the past 15 years at CERN a lot of scripts and panels have been developed for the SCADA tool. This created an important base of legacy work that cannot be deprecated. The Web UI has to be compatible with existing developments. Yet, at the same time it should remain flexible and open to improve the existing panels with user experience paradigms the web can offer.

Scalability: The new Web UI infrastructure should support up to a thousand simultaneous clients in read-only mode. This is for the use case of broadcasting a machine state to the many users interested in these information. One related requirement is that the new UI system should minimize any extra load on the production data servers running SCADA systems. At best, a decoupling and/or buffering should exist between a SCADA system and its numerous clients.

Integration: Industrial control systems are about integrating off-the-shelves components. A Web UI for it should have the same philosophy of reusing readily available web servers and services. In particular, it should integrate with web-hosting services centrally supported within an organisation running industrial control systems. As such, it emphasises the need to use compatible web standards and be platform independent. In a similar way, it

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

A NEW DATA ACQUIRING AND QUERY SYSTEM WITH ORACLE AND EPICS IN THE BEPCII*

Chunhong Wang, Luofeng Li#,
Institute of High Energy Physics, Chinese Academy of Science, Beijing 100190, China

Abstract

The old historical Oracle database in the BEPCII has been put into operation in 2006, there are some problems such as the program operation instability and EPICS PVs loss, a new data acquiring and query system with Oracle and EPICS has been developed with Eclipse and JCA. On one hand, the authors adopt the technology of the table-space and the table-partition to build a special database schema in Oracle. On another hand, based on RCP and Java, EPICS data acquiring system is developed successfully with a very friendly user interface. It's easy for users to check the status of each PV's connection, manage or maintain the system. Meanwhile, the authors also develop the system of data query, which provides many functions, including data query, data plotting, data exporting, data zooming, etc. This new system has been put into running for three years. It also can be applied to any EPICS control systems.

INTRODUCTION

BEPCII has successfully built its control system with EPICS[1]. Besides using Archiver, a historical database[2] with Oracle had developed for a long term storage. The data acquisition program developed by Python got the EPICS PVs from the EPICS IOCs. This program was scheduled by a back processing Cron. When it is scheduled timely, many CA were recreated and took too much times. As a consequence, the first schedule had been not finished, the next schedule had started. Many such data acquisition programs had been running in parallel so that the control network congestion caused EPICS PVs loss. Such kind of problem had happened many times since 2006. So, it is necessary to reconstruct a new historical database and data acquiring and query system with Oracle. Since EPICS has been widely applied to accelerator control systems, it's becoming a hot topic how to store EPICS PV data into a database for a long term storage. So, the purpose of this new system development is to be commonly used to EPICS control system.

THE OLD ORACLE SYSTEM

The old data acquiring and query system[2] with Oracle consisted of three parts: data acquisition and Oracle database and data query (see Fig. 1).

The Old Data Acquisition Program

The data acquisition software programmed by Python got the EPICS PVs from the EPICS IOCs.

It was scheduled by a Cron processing in Linux. It has the following disadvantages:

- Ineffective data acquisition sampling rate (once per minute). It took many times for the CA channels reconnection when it was executed.
- The program was repeatedly performed periodically. It means the next schedule has started when the first program had not been finished. Many same programs were running simultaneously so that the control network congestion caused PVs data loss.

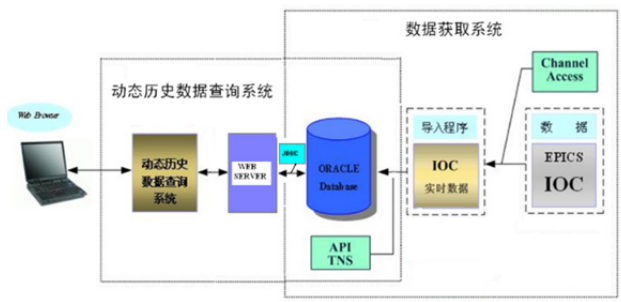


Figure 1: The old Oracle System[2].



Figure 2: the Interface of the old Query System[2].

The Old Query System

The old data query system was in B/S developed using Java. The query interface is shown in figure 2. It has the following disadvantages:

- It's impossible to implement the correlated subquery between the subsystems.
- The database tables are designed into one independent table for each subsystem (e.g. power supply control system table and vacuum control system table and so on).
- There is no any zoom in function and data without time stamp display in the interface.

*supported by NFSC(1137522)
#on leave from IHEP

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

PYTHON SCRIPTING FOR INSTRUMENT CONTROL AND ONLINE DATA TREATMENT

N. Xiong, N. Hauser, D. Mannicke, ANSTO, Sydney, Australia

Abstract

Scripting is an important feature of instrument control software. It allows scientists to execute a sequence of tasks to run complex experiments, and it makes a software developers' life easier when testing and deploying new features. Modern instrument control applications require easy to develop and reliable scripting support.

At ANSTO we provide a Python scripting interface for Gumtree [1]. Gumtree is an application that provides three features; instrument control, data treatment and visualisation for neutron scattering instruments. The scripting layer has been used to coordinate these three features. The language is simple and well documented, so scientists require minimal programming experience. The scripting engine has a web interface so that users can use a web browser to run scripts remotely.

The script interface has a *numpy*-like library that makes data treatment easier. It also has a GUI library that automatically generates control panels for scripts. The same script can be loaded in both the workbench (desktop) application and the web service application for online data treatment. In both cases a GUI will be generated with similar look and feel.

INTRODUCTION

The Gumtree scripting interface has a *numpy*-like Python library that makes data treatment easier. It also has a graphical user interface library that automatically generates control panels for scripts. The generic GUI rendering feature allows the same scripts to be loaded in both the desktop applications and the web service applications. The scripting interface benefits both the users and the developers. Users can easily make scripts to run experiments or treat the data, with a graphical interface automatically created for these scripts. Developers save a lot of time deploying new products when using this feature since they don't have to recompile. The total cost of operating this framework is therefore significantly reduced

SYSTEM ARCHITECTURE

It is always important to provide a decent scripting feature in scientific software. When users want to run an experiment, it is a more flexible and direct way to write a script to finish some complex experiment steps. A smart graphical user interface such as software wizard is very useful, but has limited scope. Once users become more familiar with the instrument and the software, they may want to write ad hoc scripts to run complex experiments. In another scenario, when users want to do data treatment on the experimental results, experienced users write their own scripts to do it. In this way they know exactly what will be done on their data and the scripts are always under their control. Many users love to write the scripts themselves. Other users who don't do programming still agree on the importance of scripting. And they will seek support from the instrument team to write scripts for them.

Gumtree is an application that provides general framework for building scientific software [2]. It is customised at ANSTO for neutron scattering experiments. It provides experiment control and data treatment services for 9 neutron scattering instruments in the Bragg Institute, ANSTO. It is designed to provide a simple and reliable scripting feature for instrument users.

Gumtree software is written in the Java programming language. But the scripting feature is provided in Python language. As shown in Figure 1, the Python engine is a core service of the application.

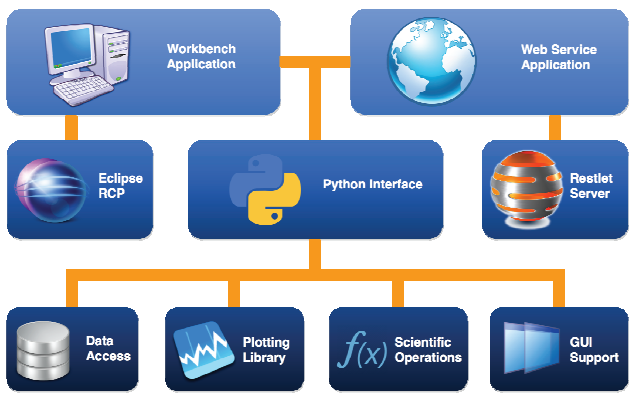


Figure 1: Gumtree system architecture.

MULTI-DIMENSIONAL ARRAY OBJECT

One of the most important scripting packages in Gumtree is the multi-dimensional array object package. It is a *numpy*-like library that provides convenient access to array operations, the Dataset interface. However it is more than just *numpy*. It provides functions for treating experiment data, such as error propagation, carrying axes information, carrying metadata, and so on. These functions make it easier for users to write scripts for data manipulation and visualisation.

Categories of the Scripting Interface

The Dataset interface borrows a number of *numpy* function names. Besides that, it also has exclusive function names to make Nexus data reduction easier.

In implementation, the python Dataset is a wrapper of a Java CDMA object [3]. The greater part of the logic

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

PARAMETERS TRACKING AND FAULT DIAGNOSIS BASE ON NoSQL DATABASE AT SSRF*

Y.B. Yan, Y.B. Leng#, L.W. Lai, Z.C. Chen
Shanghai Synchrotron Radiation Facility (SSRF)
Shanghai Institute of Applied Physics, Chinese Academy of Sciences
Shanghai 201204, P.R. China

Abstract

As a user facility, the reliability and stability are very important. Besides using high-reliability hardware, the rapid fault diagnosis, data mining and predictive analytic s are also effective ways to improve the efficiency of the accelerator. A beam data logging system was built at SSRF, which was based on NoSQL database. The logging system stores beam parameters under some predefined conditions. The details of the system will be reported in this paper.

OVERVIEW

Shanghai Synchrotron Radiation Facility (SSRF) is one of the advanced third generation light sources in the world, and it supports and pushes the cutting-edge scientific research and the innovation in China. It is composed of a 150 MeV linear accelerator, a 3.5 GeV booster, a 3.5 GeV storage ring, more than ten beamlines and experimental stations.

Due to the large amount of devices in the accelerator system, the era of big data is in fact coming with the rapid development of the related technology. Owing to the significant relationship among the data acquisition, storage and analysis, scale effect of big data has brought great challenges to the current data storage and analysis. Although the distributed data acquisition and processing system has been widely used, the distributed data storage and analysis has not been achieved due to the constraint of storage, network, and others. Some preliminary attempts have been making at SSRF, mainly based on the beam instrumentation and control system. This work would unearth more value of runtime data, which would also improve the stability and reliability of existing accelerator system.

SSRF beam instrumentation system consists of more than 200 devices, which covered the beam position measurement, beam charge & current measurement, beam size & length measurement, fill pattern measurement and so on [1]. All these parameters are very important during the accelerator commissioning, operation and machine studies. More than 20k scalar process variables and hundreds of 2k-points waveform records are published online per second. With proper storage and analysis tool-kits, these data could be invaluable. Otherwise the potential of various advanced electronics will be wasted.

On the other hand, various hardware and software failures have been recorded in the past few years, such as global orbit disturbance, random glitch or offset jump of individual position readings [2]. All these failures affected the reliability and stability of the accelerator system. There are no effective tools to analyze the reason due to lack of adequate raw data. The regular sampling rate of achieved data is about one hertz. History of broadband data such as turn-by-turn (several hundreds kilohertz) orbit data or bunch-by-bunch data are required in this case. Due to the huge size, the data are not likely to be stored periodically. A logging system was designed and developed, which stores the broadband data under some predefined conditions.

SYSTEM ARCHITECTURE

The data logging system is based on the Couchbase [3], which is an open source, distributed NoSQL database. It provides key-value or document access with low latency and high sustained throughput. The system architecture is shown in Figure 1.

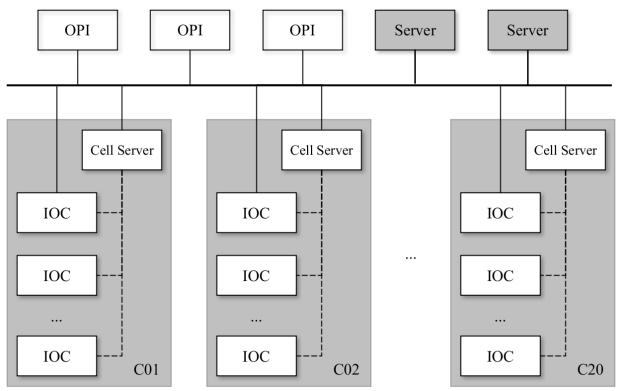


Figure 1: The architecture of data logging system.

Most part of the work is in process, especially the distributed data storage and analysis. The centralized part is mainly discussed below. In this system, IBM System x3550 M4 server and IBM Storwize V3700 storage system are adopted, which is a cost-effective option to achieve high performance. All the software run on the Linux operating system, and written using python, a widely used general-purpose, high-level programming language.

* Work supported by the Knowledge Innovation Program of Chinese Academy of Sciences
lengyongbin@sinap.ac.cn

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

FPGA FIRMWARE FRAMEWORK FOR MTCA.4 AMC MODULES*

Lukasz Butkowski, Tomasz Kozak, Bin Yang, DESY, Hamburg, Germany
Paweł Prędki, DMCS, Lodz University of Technology, Lodz, Poland
Radoslaw Rybaniec, ISE, Warsaw University of Technology, Warsaw, Poland

Abstract

Many of the modules in specific hardware architectures use the same or similar communication interfaces and IO connectors. MicroTCA (MTCA.4) is one example of such a case. All boards: communicate with the central processing unit (CPU) over PCI Express (PCIe), send data to each other using Multi-Gigabit Transceivers (MGT), use the same backplane resources and have the same Zone3 IO or FPGA mezzanine card (FMC) connectors. All those interfaces are connected and implemented in Field Programmable Gate Array (FPGA) chips. It makes possible to separate the interface logic from the application logic. This structure allows to reuse already done firmware for one application and to create new application on the same module. Also, already developed code can be reused in new boards as a library. Proper structure allows the code to be reused and makes it easy to create new firmware.

This paper will present structures of firmware framework and scripting ideas to speed up firmware development for MTCA.4 architecture. European XFEL control systems firmware, which uses the described framework, will be presented as example.

INTRODUCTION

The MTCA.4 standard is derived from the Advanced Telecommunication Computing Architecture. It is enhanced for Rear I/O and Precision Timing and offers a compact environment for transmission and parallel processing of large amounts of data. The mechanics and connectivity is defined by the standard PICMG MTCA.4 specification [1,2]. Main modules are called Advanced Mezzanine Carriers (AMC). Each of them can be paired with the Rear Transition Modules (RTM); front-to-rear communication and signal transfer is handled over the so-called Zone 3 region. The basic architecture follows the idea of a centralized powerful processing unit that is connected to various AMC I/O boards over several PCIe lanes, dedicated trigger lines, clock lines and platform related management lines. The AMC backplane offers also 4 ports for low-latency-links connections (eight differential pairs, full-duplex) that can reach up to 10 Gbit/s allowing the boards to communicate using serial transmission (e.g. utilizing the MGTs).

The main function of the AMC modules is to provide communication interfaces and perform digital signal processing. All I/O lines are connected to the FPGA chip on AMC board, shown in Figure 1. Depending on the application and the used AMC, different RTM, FMC or front I/O modules can be connected. Thanks to flexibility of the FPGAs, many different applications can be run on one AMC, still using the same hardware resources If the support of hardware is already provided in the code, it can be reused in many applications. The MTCA.4 firmware framework introduces an additional abstraction layer that separates the hardware dependent logic from user application logic. The framework specifies universal interfaces on this layer. This allows the same firmware and software components to be reused, irrespective of the type of the used hardware. It means that one application logic can be run on different AMC modules.

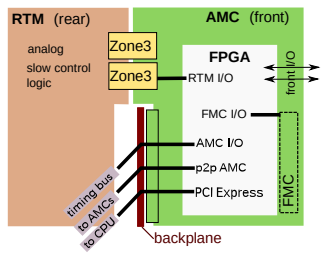


Figure 1: MTCA.4 front and rear modules

For the European XFEL all AMC boards will have one common feature: Xilinx FPGA chips, which perform, among other tasks, communication with the CPU as well as data acquisition and data processing required in control systems. All firmware algorithms are written using hardware description language (VHDL) code.

In the following sections of this paper the structure of a typical VHDL project used at European XFEL MTCA.4 [3] systems is presented.

GENERAL STRUCTURE

The code for the FPGA firmware was divided into tree main parts: board, board payload and user application.

The board section is hardware dependent and is specific for a given board. One AMC has one board support package (BSP). It has all components for configuration of peripherals on the board, such as clocks, I/O buffers, Analog to Digital Converters (ADCs), Digital to Analog Converters (DACs), etc. It is responsible for implementing communication interfaces with other boards and with the CPU. It also includes Direct Memory Access (DMA) and Data acquisition (DAQ) on DDR memory chips. The board part of the firmware provides interfaces for all resources on board.

Board payload is a middle layer section, connecting the board interfaces with the application. It contains all the

*Work supported by DESY MSK group.
#lukasz.butkowski@desy.de
ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

ENCODER INTERFACE FOR NSLS-II BEAM LINE MOTION SCANNING APPLICATIONS

Ruslan Kadyrov#, Kiman Ha, Eli Stavitski, Joseph De Long, Sung-Leung So,
BNL, Upton, New York, USA

Abstract

The NSLS-II synchrotron started operations at BNL, USA. Whereas first beamlines welcome first users, the other beamlines are under design or construction. The variety of motion control applications on existing and future NSLS-II beamlines demand custom control electronics developed to meet specific needs and ease integration into existing systems. Thus an encoder interface was designed for a number of detection techniques that require a live position capture. It implements four identical individually isolated channels with encoder input capturing and output quadrature encoder interface generated. Encoder input handles an incremental quadrature encoder with a digital RS-422A interface and output frequencies up to 10 MHz. The logic, based on Xilinx Virtex-6 FPGA, processes signals from an encoder and associates it with the machine timestamp. Several filtering and compression techniques are also applied. The device then retranslates the interface signals for the motion controller, allowing the device to be installed between encoder and motion controller with no interference to the system. Captured data is streaming to the server using TCP/IP stack, with the server side running an EPICS IOC. The device status and control registers values are also available through EPICS channel access. All operator screens are developed with Control System Studio toolset. The design harmoniously complements the family of the NSLS-II equipment sharing same mechanical and electrical platform. This paper describes the functional design, configuration, operation and current status of the encoder interface design.

INTRODUCTION

Motion scanning applications role in X-Ray experiments consists in cross correlation analysis with data obtained from various optical equipment. Position is usually detected by a quadrature incremental or absolute encoder. Delta Tau GeoBrick motion controller used in NSLS-II motor-based applications doesn't store or timestamp reported position. To keep the design simple, same encoder was decided to be used both for the axis motion control and position capturing. Thus the design should be able to retransmit the encoder signals for the motion controller without interference to the work of the latter.

The logic design concept was inspired by a similar development at Diamond Light Source Ltd [1]. The design instrumentation base shares existing hardware with several devices used for the NSLS-II accelerator FPGA-based applications: power supply controllers [2], beam position monitors [3], cell controller.

SYSTEM REQUIREMENTS

The design should provide applicability for various motion scanning applications, easy integration into existing machine control system, and compliance with the NSLS-II beamline instrumentation requirements [4]:

- The design should be able to handle a Renishaw Quadrature incremental encoder with 10 MHz operating frequency aiming to retransmit encoder interface signals for a motion controller allowing bypass connection with minimum or no interference to the encoder connection,
- Relative encoder position is to be timestamped and send to a server allowing to restore an encoder position with 10 kHz sampling rate,
- Standard EPICS driver should be used to minimize device support development needs,
- Ability to start capturing on external trigger or on EPICS PV write,
- Compatibility with existing timing system,
- Flexibility to custom experiment needs, e.g. generating triggers for external devices at certain axis positions or every N steps.

DESIGN COMPONENTS

Instrumentation Base

The design fits in a 2U 19" chassis with hardware based on standard NSLS-II platform. It is a metal enclosure with two power supplies, minimum internal wiring and place for two standard boards mounted on the bottom plate using stand-offs. Front-face mounted FPGA-based main board is in charge of purely digital part of the design, and secondary rear-face mounted interface board for device-specific circuitry. This modular design makes it possible to change functionality or build another device by putting a new interface board and developing a custom firmware for FPGA.

Hardware

The hardware leverages the NSLS-II BPM Digital Front End (DFE) board with Virtex-6 FPGA and periphery [3]. A custom interface board was developed to handle encoder interface specific circuitry. Four identical channels are implemented with each capable to receive and form the quadrature encoder interface signals.

Input and output interfaces are connected to double stacked D-sub 15-pin connectors along the rear panel of

rkadyrov@bnl.gov

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

Em# PLATFORM: TOWARDS A HARDWARE INTERFACE STANDARDIZATION SCHEME

O. Matilla, J. Avila-Abellan, M. Broseta, G. Cuní, D. Fernández-Carreiras, A. Ruz, J. Salabert, X. Serra-Gallifa, ALBA-CELLS Synchrotron, Cerdanyola del Vallès, Spain

Abstract

Low current measurements developments have been pointed out as strategic for ALBA synchrotron. From the first day of operation of the seven Beamlines currently in operation ALBA Em electrometer has been successfully used. Today, the two new beamlines of Phase 2 that are being constructed and the new end stations have required specification changes in terms of increased accuracy, capability of synchronization, timestamping, management of large buffers and high performance closed-loop implementation. The scheme of full custom hardware design has been abandoned. ALBA Em# project approach has been based in the selection of industry standard interfaces: FMC boards design for custom needs, FMC carrier over PCIe using SPEC board from CERN under OHWR license, and Single Board Computer using PCIe to implement interfaces with the control system. This Paper describes the new design of the Electrometers at Alba, suited for the newer requirements, more flexible, performing and maintainable, which profits from the know-how acquired with previous versions and suits the new data acquisition paradigm emerged with the standardization of quick continuous scans and data acquisition.

INTRODUCTION

ALBA is a 3GeV third generation synchrotron light source located in Cerdanyola del Vallès (Barcelona). ALBA houses 7 commissioned beamlines and 2 more are being constructed in the facility construction Phase 2.

Low current measurements at ALBA for diagnostic or experiment applications have been performed thanks to the successful development of a 4-channel electrometer (ALBA Em) [1]. The project started in 2011 and more than 40 units are currently working at ALBA facility since then. Due to the increase of needs mainly in the experimental area an equipment redesign need raised in 2013 (Em#) [2]. Since then the project philosophy has evolved from the full custom hardware design to be based, as much as possible, in commercial components using expected long-life interfaces. The inclusion in the equipment of a Single Board Computer (SBC) integrating part of the control system increases the complexity of the project and positively forces the incorporation of software developers in the project development stage.

With the common objective of fostering the development of Em# project a collaboration agreement has been signed between ALBA and MAX IV. The agreement could be open to other institutes or facilities interested in the project.

HARDWARE APPROACH: DON'T SELECT THE COMPONENTS, SELECT THE INTERFACES

One of the lessons learned from first ALBA Em is that the obsolescence of commercial modules (like microprocessor cores) is a crucial factor in the equipment design. If an electronic module gets obsolete and its interface is not standard, the module replacement implies a dramatic effort of hardware redesign.

The project new approach is based first on selecting the industry standard interfaces intended to be used between the modules included in the equipment, and second selecting the modules that could fulfil these requirements. In this way, if a module gets obsolete, it can be replaced by another one using the same interfaces, and the redesign effort is lower.

In the Em# project three standard interfaces have been selected:

1. Ethernet: External communication interface from the embedded SBC to the Control System.
2. Peripheral Component Interconnect Express (PCIe): Internal interface between the SBC and the FPGA module.
3. FPGA Mezzanine Card (FMC): Interface between FPGA module and custom electronics based on VITA 57 standard [3].

The block diagram of the new hardware approach (see Fig. 1) is composed mainly by 3 well-bounded blocks, where the connections between them are the cited selected interfaces. The first block is a selected commercial fanless SBC including Ethernet and PCIe interfaces with a small 100mm x 100mm long life form factor: Intel NUC (DE3815TYBE). The second block is a FPGA module including PCIe and FMC (as carrier) interfaces: SPEC, a FPGA board developed by CERN and available under Open Hardware License (OHL) [4]. The third block contains custom electronics developed for the project. Custom electronics block includes an FMC board developed at ALBA under OHL license [5]. Its core is an isolated 4-channel 400kS/s 18 bits ADC capable of operating under ground voltage bias condition. This ADC FMC board has also a custom digital interface to control the rest of electronics; for the equipment I/O and the 4 ALBA Current Amplifiers. The 4 ALBA Current Amplifiers and the ADC of the FMC are designed to work with biased voltage ground, so the isolation is a key point with many implications in the hardware design.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

SINGLE NEUTRON COUNTING USING CCD AND CMOS CAMERAS

P. Mutti*, E. Ruiz-Martinez, M. Platz, P. Van Esch
Institut Laue-Langevin, Grenoble, France
M. Crisanti, Universitá degli studi di Perugia, Perugia, Italy

Abstract

The present paper explores the technical possibility to use classical imaging detectors such as Charge-Coupled Devices (CCD) and Complementary Metal-Oxide Semiconductor (CMOS) cameras as counting detectors at lower counting rates and transits smoothly to continuous imaging at higher rates. The gamma-rejection capabilities as well as the counting efficiency have been investigated with both an AmBe neutron source and in-beam experiments.

INTRODUCTION

Neutrons have mass but no electrical charge. Because of this they cannot directly produce ionisation in a detector, and therefore cannot be directly detected. This means that neutron detectors must rely upon a conversion process where an incident neutron interacts with a nucleus to produce a secondary charged particle. Classical neutron detectors are based upon particle detection technologies like gas filled proportional counters or scintillation detections. These detectors have a high dynamic range, and are very performing at low counting rates and fast timing (time-of-flight) applications. At high counting rates however, continuous imaging detectors such as CCD or CMOS camera's optically linked to scintillators, can have very good performances concerning linearity and spatial resolution but the dynamic range of these systems is limited by electronic noise and gamma background.

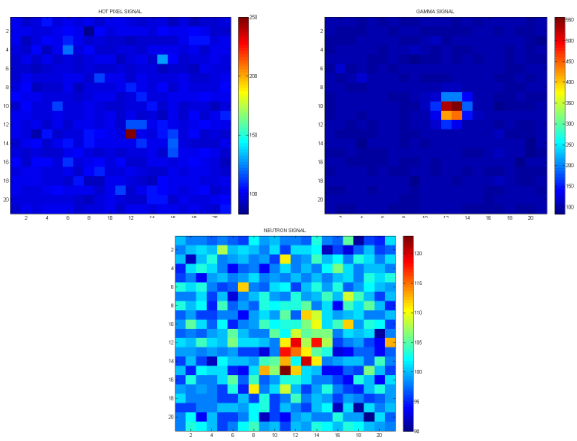


Figure 1: Different light patterns in a CMOS camera.

The capability of single neutron counting relies on the possibility to distinguish the neutron signature on the camera from those belonging to gammas or electronic noise (hot pixel). Neutron detection involves reactions releasing energies of the order of the MeV, while X-ray detection releases energies of the order of the photon energy ($\sim$10 keV). This 100-fold higher energy released in the interaction, transforms into a much different light yield allowing the individual neutron light signal to be significantly above the noise level and therefore opening the possibility for gamma-discrimination and individual counting. Figure 1 shows the different light patters obtained using our CMOS camera. The *hot pixel* is shown in the left-upper plot, the gamma detection in the right-upper and the neutron signature in the lower one.

NEUTRON RECOGNITION

^{3}He based proportional gas neutron detectors have an impressive gamma rejection ratio in the range from 10^{-6} to 10^{-8} [1]. This impressive noise rejection ratio comes from a non-linear operation, called discrimination. The noisy analogue signal is put to 0 in the end result, as long as it doesnt cross a threshold. In such a way, all the integrated noise (whether electronic noise, inherent detector noise, or gamma radiation) is eliminated entirely. If one assumes a Gaussian noise with a standard deviation σ and a threshold set at 6σ, the probability per independent sample (about every μs) to cross that positive threshold equals about 10^{-9} which gives a background rate of 10^{-3} counts/s. At a count rate of 10 neutrons per second, an integrating detector summing all the samples over a period of 1 second would have a total signal of about 60σ (assuming the individual signal level equal to 6σ) and a noise level equal to 1000σ: in other words, the signal would be entirely swamped by the noise. When used in classical integrating imaging mode, a CMOS like camera suffers a similar degradation of the final signal-to-noise ratio and, therefore, it can only really be used when the neutron fluxes are very intense, such as in neutrography, or in Laue diffraction. Normally, the camera pixel size is smaller than the size of the typical light spot from a neutron scintillator event on the camera and, therefore, one can assume that the scintillating light has a 2-dimensional Gaussian distribution as shown in Fig. 2.

We will consider a re-pixelisation where we combine the camera pixels into what we call *spot pixels*: pixels that are about the size of the FWHM of this Gaussian distribution. We take images with such short exposure times that the amount of noisy photo-electrons in a spot pixel is well below the expected number of photo- electrons corresponding to a neutron scintillation event, and we put a threshold that distinguishes them. The first problem related to this

* mutti@ill.eu

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

NEW DIGITISERS FOR POSITION SENSITIVE ^{3}He PROPORTIONAL COUNTERS

P. Mutti*, E. Ruiz-Martinez, K. Ollivier, M. Platz, P. Van Esch
Institut Laue-Langevin, Grenoble, France

Abstract

The present paper reports on the development of a new digital front-end electronics for position sensitive ^{3}He neutron detectors. The classical analog approach for the charge readout, consisting of a shaping amplifier coupled with a peak sensing ADC, has been replaced by a 64 channels, 62.5 Msample/s and 12 bit digitiser. Excellent results have been obtained in terms of position resolution and signal to noise ratio when adopting a continuous digital filtering and gaussian shaping.

INTRODUCTION

^{3}He gas-filled detectors are a classical choice for the detection of thermal and cold neutrons. The incident neutrons are captured by the ^{3}He producing a tritium and an hydrogen which are sharing the 756 keV of energy generated in the reaction.

$$^3He + n \rightarrow {}^3H + {}^1H + 756keV \quad (1)$$

The electron avalanche initiated by the 2 ions generates the detector signal.

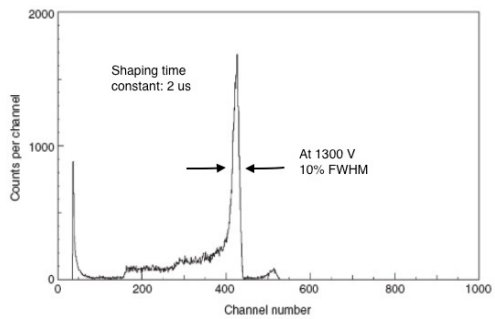


Figure 1: Pulse-height spectrum from ^{3}He tube.

Figure 1 shows a typical pulse-height spectrum from a ^{3}He tube at low gain. At higher gain, non-linear effects set in and the proportionality is lost. This limits the charge one can reasonably obtain from a neutron capture event to about 1 pC. The shape of the spectrum is due to the kinematic of the reaction but also to the choice of the amplifier time constant. The full peak is originated by the collection of the total energy of the two ions (765 keV). If one of the ions is absorbed by the tube walls then, the total collected energy is smaller. This results in the peak asymmetry visible at the left part of the peak in Fig. 1. A choice of a shaping time ranging from 0.5 to 2 μs is a good compromise between good resolution and high count-rate capability.

* mutti@ill.eu

The classical geometry of a charge-division neutron detector consists of a cylindrical volume housing a resistive anode (see Fig. 2).

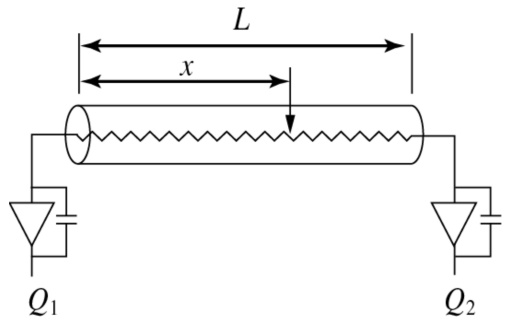


Figure 2: Principle of a charge-division neutron detector.

Electrical signals are extracted at both ends of the tube and the information about the interaction point along the tube can be derived by the ratio of the collected charged at both ends. The theoretical spacial resolution achievable for a detector of total length L is function of the temperature T of the resistive electrode, of the total detector capacity C and of the collected charge. This can be expressed by

$$\frac{FWHM}{L} = \frac{2.54 \cdot \sqrt{kTC}}{Q_1 + Q_2} \quad (2)$$

When using low-noise amplifiers, one can expect from the previous equation (2) a spacial resolution in the order of 0.1% of the tube length. However, Eq. (2) is not applicable in the case of ^{3}He neutron detectors because the charge collection time is much longer than the RC time constant of the detector, as shown in Fig. 3. This results in an actual spacial resolution of the order of 1% of the tube length. The limiting factor being dominated by the noise of the resistive wire and by the integration time.

FRONT-END ELECTRONICS

Our current front-end analog electronics for charge-division consists of a shaping amplifier containing a 4th order Gaussian filter, a baseline correction circuit and a single-shot 12 bit ADC [1], while the position is coded with a 8-bit word due to the limited resources available. The essential idea is to implement digitally an equivalent signal treatment using the CAEN type V1740 [2] digitiser board with an ad-hoc firmware. Since the pulse-shape of the signals depends on the interaction location and on the charges propagation direction (see Fig. 3), the conventional pulse-shape analysis is inappropriate.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

THE CONSTRUCTION OF THE SUPERKEKB MAGNET CONTROL SYSTEM

T. T. Nakamura#, A. Akiyama, M. Iwasaki, H. Kaji, J.-I. Odagiri, S. Sasaki, KEK, Ibaraki, Japan
N. Yoshifuji, EJIT, Hitachi, Ibaraki, Japan
T. Aoyama, T. Nakamura, K. Yoshii, Mitsubishi Electric System & Service Co., Ltd, Tsukuba, Japan

Abstract

There were more than 2500 magnet power supplies for KEKB storage rings and injection beam transport lines. For the remote control of such a large number of power supplies, we have developed the Power Supply Interface Controller Module (PSICM), which is plugged into each power supply. It has a microprocessor, ARCNET interface, trigger signal input interface, and parallel interface to the power supply. The PSICM is not only an interface card but also controls synchronous operation of the multiple power supplies with an arbitrary tracking curve. For SuperKEKB we have developed the upgraded version of the PSICM. It has the fully backward compatible interface to the power supply. The enhanced features includes high speed ARCNET communication and redundant trigger signals. Towards the Phase 1 commissioning of SuperKEKB, the construction of the magnet control system is ongoing. First mass production of 1000 PSICMs has been completed and their installation for Phase 1 configuration has been done. The construction status of the magnet control system is presented in this report.

INTRODUCTION

KEKB is the asymmetric energy electron-positron collider, which is dedicated to the B-meson physics. Its operation started in December 1998 and finished in June 2010. The KEKB accelerator control system has been constructed based on EPICS (Experimental Physics and Industrial Control System) tool kit. EPICS provides core mechanism for the distributed control system. EPICS runtime database is running on a local control computer called IOC (Input/Output Controller). More than 100 VME/VxWorks computers have been installed as IOC in the KEKB accelerator control system.

SuperKEKB is the upgrade of KEKB. It aims at 40-times higher luminosity than the world record by KEKB. Its operation is going to be performed in three phases. The Phase 1 is the operation without QCS (final focusing superconducting magnets) nor BelleII detector. Start of the Phase 1 is scheduled in Feb. 2016. Phase 2 follows after installation of QCS and BelleII,.

In the KEKB storage rings and the injection beam transport lines, about 2500 magnet power supplies were installed [1] and controlled by 11 IOCs. To connect such a large number of power supplies to the IOCs, we adopted ARCNET as the field bus and developed the PSICM (Power Supply Interface Controller Module) [2]. The PSICM is the ARCNET interface board for the power supply. The PSICM has the shape of 3U Euro-card format (100mm × 160mm) with a DIN 64-pin connector and can be plugged into the power supply. Figure 1 shows the photo picture of the PSICM and the magnet power supply with PSICM. Figure 2 shows the details of the interface between the PSICM and the power supply.

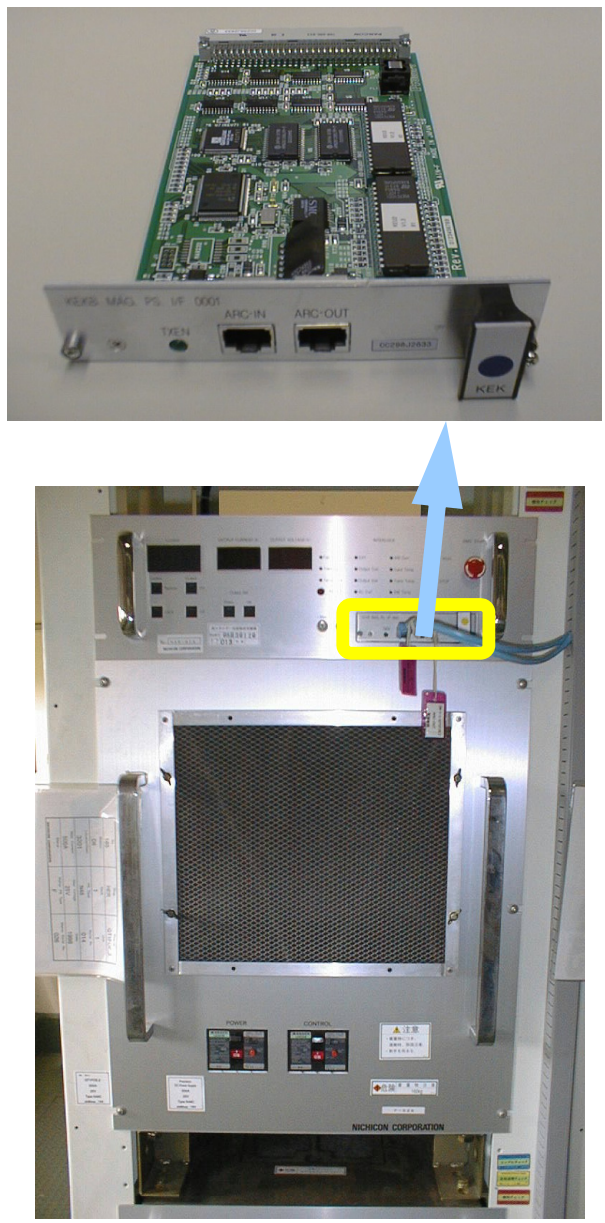


Figure 1: Whole shape of the PSICM (above) and the PSICM plugged in a magnet power supply (below).

tatsuro.nakamura@kek.jp

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

CERN OPEN HARDWARE EXPERIENCE: UPGRADING THE DIAMOND FAST ARCHIVER

I.S. Uzun and M. Abbott, Diamond Light Source, Oxfordshire, UK

Abstract

Diamond Light Source developed and integrated the Fast Archiver [1] into its Fast Orbit Feedback (FOFB) communication network in 2009. It enabled synchronous capture and archive of the entire position data in real-time from all Electron Beam Position Monitors (BPMs) and X-Ray BPMs. The FA Archiver solution has also been adopted by SOLEIL and ESRF. However, the obsolescence of the existing PCI Express (PCIe) based FPGA board from Xilinx and continuing interest from community forced us to look for a new hardware platform while keeping the back compatibility with the existing Linux kernel driver and application software. This paper reports our experience with using the PCIe SPEC card from CERN Open Hardware initiative as the new FA Archiver platform. Implementation of the SPEC-based FA Archiver has been successfully completed and recently deployed at ALBA in Spain.

INTRODUCTION

The Diamond FOFB system is operating successfully and achieving the required level of suppression of beam disturbances since 2007 [2]. The system architecture for the FOFB for a single Storage Ring (SR) cell is shown in Fig. 1. Diamond SR consists of 24 cells, and overall FOFB architecture consists of:

- 172 Electronic Beam Position Monitors (BPMs) located around the storage ring.
- 24 VME-based processors, each populated with a PMC receiver, running the feedback algorithm.
- Power supply controller units for 168 slow corrector magnets.
- Communication network to distribute the beam position values.

BPMs are used for electron beam diagnostics. In the BPM FPGA the sampled button intensities are converted to X,Y positions and reduced by decimation and filtering to machine revolution frequency (534 kHz) and then to the "Fast Acquisition" (FA) data rate of 10 kHz and the "Slow Acquisition" rate of 10 Hz for EPICS clients. The FA data is broadcast by Diamond Communication Controller IP (DCC) [3] on the FPGA through a dedicated communication network, which makes this data stream simultaneously available to all PMC receiver nodes on VME processors for fast feedback processing.

Following the successful operation of FOFB system, a need for recording longer periods of synchronous position data from all nodes on the communication network (the whole orbit) at FA data rate was identified. The DCC IP and dedicated communication network enabled us to additional nodes with little or no impact on the operation of the FOFB; and the FA Archiver is one such node implemented on a PCIe hardware [1].

Two FA Archiver nodes have been running at Diamond, one for Booster and one for Storage Ring since 2009, giving us access to records of up to two weeks of position data across the both rings. We have been using this data for:

- Identification and detailed investigation of any orbit events that happened in the recent past.
- Correlation of orbit events with other measurements by adding X-ray BPMs and RF cavity voltage readings to the FA network.
- Model analysis of orbit motion over longer periods for the optimisation of the fast orbit feedback.

FA Archiver interface was implemented on a commercial Xilinx ML555 PCIe FPGA [4] board as a cost-effective solution so that we could focus our efforts on firmware development and software integration. Following its successful commissioning at Diamond, SOLEIL and ESRF integrated FA Archiver as part of their fast feedback system. This was possible since both facilities built their FOFB data distribution network using DCC giving them the protocol compatibility required.

Recent interest from ALBA as the latest user of DCC, and Brazilian Light Source (LNLS) for evaluation purposes brought the obsolescence of the ML555 platform to our attention. In order to support the community, we took on an upgrade project targeting implementation of FA Archiver on a modern platform with longer term support and availability. Following evaluation of commercially available products and in-house built possibilities, we have decided that SPEC card from CERN Open Hardware initiative would be the best

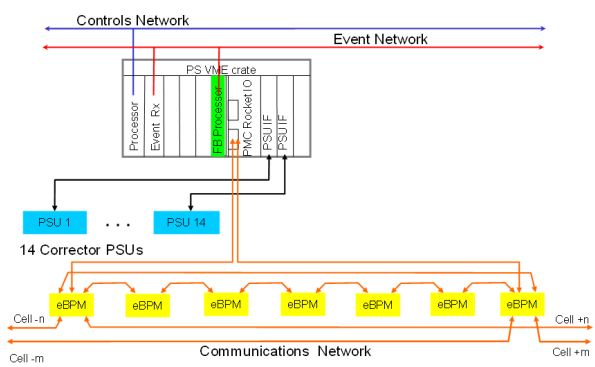


Figure 1: FOFB architecture for single SR Cell. It is repeated for each 24 cells in Diamond SR.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

DESIGN OF EPICS IOC BASED ON RAIN1000Z1 ZYNQ MODULE*

Tao Xue, Hongming Li, Guanghua Gong, Jianmin Li, Tsinghua University, Beijing, China

Abstract

ZYNQ is the new architecture of FPGA with dual high performance ARM Cortex-A9 processors from Xilinx. A new module with Giga Bit Ethernet interface based on the ZYNQ XC7Z010 is development for the High Purity Germanium Detectors' data acquisition in the CJPL (China JinPing under-ground Lab) experiment, which is named as RAIN1000Z1. Base on the nice RAIN1000Z1 hardware platform, EPICS is porting on the ARM Cortex-A9 processor with embedded Linux and an Input Output Controller is implemented on the RAIN1000Z1 module. Due to the combine of processor and logic and new silicon technology of ZYNQ, embedded Linux with TCP/IP sockets and real time high throughput logic based on VHDL are running in a single chip with small module hardware size, lower power and higher performance. We will introduce how to porting the EPICS IOC application on the ZYNQ based on embedded Linux and give a demo of IO control and RS232 communication.

INTRODUCTION OF RAIN1000Z1

ZYNQ [1] is developed by Xilinx for high performance real time processing and embedded application. In the chip of ZYNQ, dual core of ARM Cortex-A9 processor with rich peripherals is the half of the chip, the other half part is the FPGA. People can develop the embedded application in the processor with C programming, such as embedded Linux and uC/OS or VxWorks real time OS can be running on the Cortxe-A9 dual core processors. In the other part of FPGA, people can develop the real time processing logic or DSP function with VHDL/Verilog programming. It is amazing that combine the two parts in one chip, the total footprint is more compact, board space is saved and the power is saved. At the same time, the assistive circuit, such as the on board power circuit, supervisor circuit and debug circuit is simpler.

With this feature of ZYNQ, we develop a readout module for the High Purity Germanium Detectors' data acquisition in the CJPL (China JinPing under-ground Lab) experiment, which is named as RAIN1000Z1 [2]. Figure 1 shows the system architecture of the RAIN1000Z1 module.

The RAIN1000Z1 module is based on the XC7Z010 chip of ZYNQ series, in the CLG225 footprint, which is the smallest package of ZYNQ series. A 256Mbytes DDR3 SDRAM running at 1066MHz and 1000Mbps Ethernet PHY chip 88E1512 from Marvell Inc. are embedded in the RAIN1000Z1 module. Also the power circuit and Flash for firmware saving in QSPI interface are all integrated in the module. Figure 2 is the top view of the RAIN1000Z1 module.

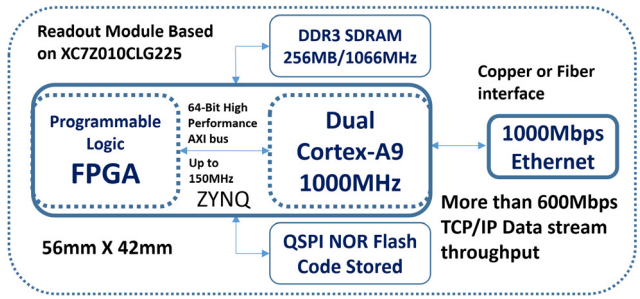


Figure 1: Architecture of the RAIN1000Z1 module.

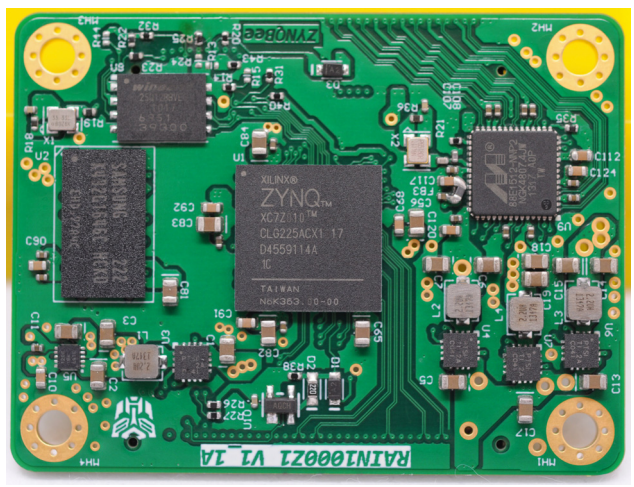


Figure 2: Top view of the RAIN1000Z1 module.

China JinPing underground Laboratory (CJPL) is a deepest underground in the world and provides a very good environment for direct observation of dark matter.

The CDEX experiment will going to direct detect the interaction of WIMP with nucleon in CJPL with high sensitivity in low mass region. Both CJPL and CDEX have got much more progress in recent two years. The CDEX use point contact germanium semi-conduct detector PCGe which has less than 300eV threshold.

The RAIN1000Z1 is used in the readout for HPGe detectors for the CDEX (China Dark Matter Experiment) in CJPL, we can easily get more than 600Mbps TCP/IP data throughput [2].

In the RAIN1000Z1 module, the embedded Linux is running. Currently, we use the version of 3.19 from Xilinx support git hub. Figure 3 shows the booting log of embedded Linux.

Some GPIO and one UARTlite (use the IP core from Xilinx in the Vivado development environment) are im-

*The authors are with Key Laboratory of Particle & Radiation Imaging, Department of Engineering Physics, Tsinghua University, Beijing, China. 100084 (email: gbe.tao.xue@gmail.com).

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

A FORMAL SPECIFICATION METHOD FOR PLC-BASED APPLICATIONS

D. Darvas*, E. Blanco Viñuela, CERN, Geneva, Switzerland
I. Majzik, Budapest University of Technology and Economics, Budapest, Hungary

Abstract

The correctness of the software used in control systems has been always a high priority, as a failure can cause serious expenses, injuries or loss of reputation. To improve the quality of these applications, various development and verification methods exist. All of them necessitate a deep understanding of the requirements which can be achieved by a well-adapted formal specification method. In this paper we introduce a state machine and data-flow-based formal specification method tailored to PLC modules. This paper presents the practical benefits and new possibilities of this method, comprising consistency checking, PLC code generation, and checking equivalence between the specification and its previous versions or legacy code. The usage of these techniques can improve the level of understanding of the requirements and increase the confidence in the correctness of the implementation. Furthermore, they can help to apply formal verification techniques by providing formalised requirements.

INTRODUCTION AND MOTIVATION

The complexity of process control systems is steadily increasing, resulting in more and more complex control software. Without an appropriate specification it is increasingly difficult to understand the requirements, hence to develop and maintain the programs.

The motivation of this work originates from CERN (the European Organization for Nuclear Research), where numerous control systems are in use to operate the research facilities. Many of these control systems rely on Programmable Logic Controllers (PLCs). The increasing complexity of the control systems results implies an increasing complexity issue. The non-formal, ad-hoc specification methods in use are less and less able to cope with the complexity of the systems. Therefore we propose a new specification method that aims to handle the increasing complexity of PLC program units.

Obviously, this is not the first work on specification methods for PLCs. *Grafcet* is a standardised specification method [1] based on safe Petri nets. It is convenient to describe finite state machines, but they are not universal enough to be generally used in the development of CERN's control software as the finite state machines are just one part of the final code deployed. *ProcGraph* [2] is one of the recent attempts to improve specifications of PLC programs, but it is on a low abstraction level, too close to the real program code. Other PLC-related specification methods also exist, but none of them seemed to be able to cope with the size and complexity of the PLC programs used at CERN.

Also, there are numerous general-purpose modelling methods that can be used for specifying PLC programs. One of them is the widely-known Simulink Stateflow, that has a complex semantics not adjusted to the needs of the PLC domain, making its usage difficult and potentially error-prone.

The structure of the paper is the following. The next section introduces the main concepts and the structure of the proposed formal specification method. After, the potential benefits of the new method are discussed. Finally, the last section concludes the paper and draws up the future work.

Due to space limitations, the discussion of the syntax and semantics of the specification method is introductory. The reader can find more details in our technical report [3].

MAIN SPECIFICATION CONCEPTS

The goal of our work is to provide a specification method that is a *complete*, *formal* behaviour description of specific functionality which corresponds to one or some PLC components. Therefore it can be used to describe individual PLC components with high precision. Then, these precise descriptions can be naturally composed into complete control systems in the future.

Previous work [4] discussed the requirements towards such a formal PLC specification method. The main requirements can be summarised as follows: (1) providing multiple formalisms adapted to the needs and knowledge of the PLC community, (2) keeping the "core logic" clean by decoupling the input/output-handling, (3) supporting events with proper semantics, and (4) limiting the set of possible errors by limiting the expressivity. The domain-specific requirements are complemented by general requirements, most notably the need for a formal, precise specification formalism, that is also lightweight (can be introduced with small effort) and complete (it describes all required and allowed behaviours). These requirements resulted in PLCspecif, a complete formal specification method for PLC programs.

Structure of the Specification

The main building block of the specification is the *module* that is either a composite module (its behaviour is described by some submodules) or a leaf module (its behaviour is directly described). To provide a specification method that is familiar for the PLC developers, the leaf module descriptions are based on three widely-known formalisms: state machines, data-flow diagrams, and standard PLC timers.

Each module (both composite and leaf modules) is further decomposed into three main parts: (1) input definitions, (2)

* Corresponding author. E-mail: ddarvas@cern.ch, darvas@mit.bme.hu

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

PLCverif: A TOOL TO VERIFY PLC PROGRAMS BASED ON MODEL CHECKING TECHNIQUES

D. Darvas*, B. Fernández Adiego, E. Blanco Viñuela, CERN, Geneva, Switzerland

Abstract

Model checking is a promising formal verification method to complement testing in order to improve the quality of PLC programs. However, its application typically needs deep expertise in formal methods. To overcome this problem, we introduce PLCverif, a tool that builds on our verification methodology and hides all the formal verification-related difficulties from the user, including model construction, model reduction and requirement formalisation. The goal of this tool is to make model checking accessible to the developers of the PLC programs. Currently, PLCverif supports the verification of PLC code written in ST (Structured Text), but it is open to other languages defined in IEC 61131-3. The tool can be easily extended by adding new model checkers.

INTRODUCTION AND MOTIVATION

Operating an accelerator complex to provide facilities for particle physics research involves numerous process control tasks. Many of them (such as cooling and ventilation, cryogenics, gas systems) are controlled by Programmable Logic Controllers (PLCs) at CERN, the European Nuclear Research Organization. As these systems are critical for the operation of CERN, the correctness of the executed PLC applications is a high priority.

Testing is a widely used solution to find potential failures in software. However, testing is not an universal solution for the verification of programs, for the following main reasons:

- Testing cannot show the absence of bugs, it can only show their presence.
- Testing can only check the outputs given by the software under test for some *selected input sequences* (test inputs). It cannot check efficiently general statements (e.g. "If output `FireAlarm` is true, then the output `NoAlarmPresent` should always be false.") or liveness properties (e.g. "If a request is received, a response will be sent *eventually*.").

Model checking is a good candidate to *complement* testing in order to reduce these weaknesses [1]. This paper introduces the high-level concepts of model checking and our proposed solution to incorporate model checking in the PLC software development process.

CHALLENGES OF MODEL CHECKING

Model checking is a formal verification technique that takes (1) a mathematical *model* of the system to be checked and (2) a formalized *requirement*. The model checker algorithms can decide if the given requirement is satisfied for the given model or not. Contrarily to testing, model checking checks *all possible* executions of the program and reports if any of them violates the requirement. Model checking is also able to generate *counterexamples*, i.e. input sequences demonstrating the violation of the given requirements.

However, model checking is not a silver bullet. There are three main obstacles of using this technique:

1. The model checker tools need a mathematical representation of the program. Constructing them needs lots of effort and experience in the formal methods domain.
2. The requirements should also be formalized for model checking. This is a similarly challenging task.
3. Model checking is computation- and memory-intensive. The generated models of the programs are often too large or too complex to be analysed with the available computation capacity.

These obstacles are difficult to overcome. The available tools require deep expertise in the formal verification domain. This could be the main reason why model checking is not widely used in industry yet, apart from some highly safety-critical applications in avionics, railway industry, etc.

Our goal is to provide a model checking solution for the PLC domain by overcoming the mentioned obstacles. All of them contain both theoretical and technical challenges. In earlier work [2, 3] we have provided solutions for the theoretical obstacles. All these solutions have been incorporated in a tool called PLCverif[1] that makes model checking accessible for the developers in the PLC domain. This paper focuses on this tool and on the bridge between the formal methods and PLC domains.

PLCverif: A BRIDGE BETWEEN FORMAL VERIFICATION AND PLC DOMAINS

In this section we overview the main features of the PLCverif tool. We focus on the user's point of view, therefore the structure of this section follows the normal workflow of a user.

Typical Workflow

The typical user workflow of the PLCverif tool consists of four steps:

1. Defining (importing or writing) the PLC code to be checked,
2. Defining the requirement to be verified,
3. Executing the verification,
4. Evaluating the results of the verification.

In the following sections, each step is described in detail.

* E-mail of the corresponding author: daniel.darvas@cern.ch

[1] `http://cern.ch/plcverif/`

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

CXv4, A MODULAR CONTROL SYSTEM

Dmitry Bolkhovityanov, Pavel Cheblakov, Fedor Emanov,
Budker Institute of Nuclear Physics, SB RAS, Novosibirsk, Russia

Abstract

CX control system is used at VEPP-5 and several other BINP facilities. CX version 4 is designed to provide more flexibility and enable interoperability with other control systems. In addition to device drivers, most of its components are implemented in a modular fashion, including data access at both client and server sides. The server itself is a library. This approach allows clients to access several different control systems simultaneously and natively (without any gateways). CXv4 servers are able to provide data access to clients from diverse CS architectures/protocols, subject to appropriate network module being loaded. The server library, coupled with "null link" client-server access module, allows to create standalone monolythic programs for specific small applications (such as test benches and device test screens/utilities) using the same ready code from large-scale control system but without its complexity.

CXv4 design principles and solutions are discussed and first deployment results are presented.

CX

CX was designed at BINP in late 1990s as a general control system framework which runs on Linux and several flavors of *NIX. Initially made for CAMAC, CX now supports a wide range of hardware attached via different fieldbuses, including PCI/CompactPCI, VME, CANbus, RS232/485 and Ethernet.

CX is used at VEPP-5 Injection Complex [1] and at several other BINP facilities and small-scale experiments [2–4].

VEPP-5 will supply electrons and positrons to two BINP machines — VEPP-4 [5] and VEPP-2000 [6] in the near future. VEPP-4 uses a mix of an inhouse-designed software (with roots dating back to 1970s) [7] and EPICS; VEPP-2000 employs a custom software named VCAS [8]. For VEPP-5 to carry out its mission, its control system must be able to communicate with partners' control systems.

THE PROBLEM OF INTEROPERATION

CX, like most control system frameworks used in high energy physics experiments, is distributed and is based on a 3-layer model (Fig. 1a). A typical framework is closely tied to some network protocol and is often designed jointly with a dedicated protocol. I.e., the client library implements the client side of a protocol, and server implements its side (Fig. 1b); CX belongs to this class.

However, a need to interact with a different control system framework arises sooner or later in real world. This is often solved via "gateways" — dedicated software, translating one protocol to another (Fig. 2). This approach has its disadvantages besides a necessity for an additional software component. Not only does such gateway have to convert between protocols, but sometimes between different data paradigms, it convering in a very inconvenient place.

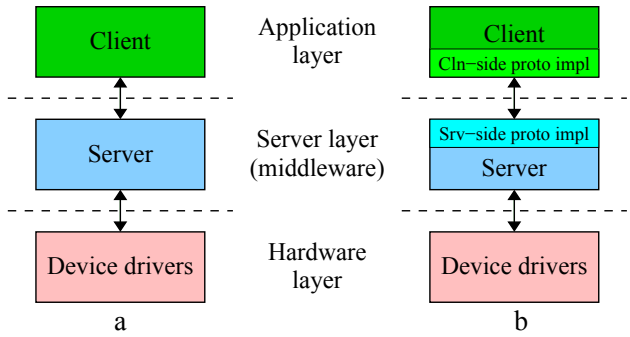


Figure 1: 3-layer architecture. (a) General layout; (b) Client-server communication.

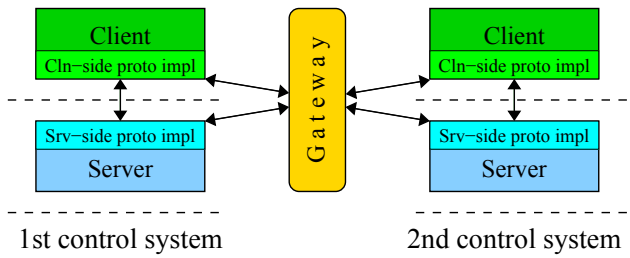


Figure 2: Interaction between different control systems via a gateway.

SOLUTION

In authors' opinion [9, p.3], the key to solution is separation of network protocol specifics and implementation from both client libraries and server. That was done in CX version 4 (see Fig. 3).

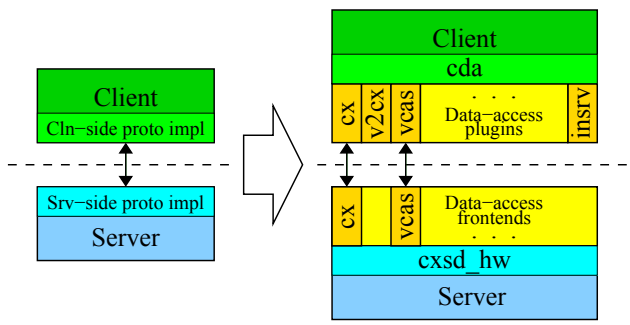


Figure 3: CXv4 modular structure.

Client library doesn't interact with server directly but rather via plugins. The library core provides clients with

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

A MODULAR APPROACH TO DEVELOP STANDARDIZED HVAC CONTROL SYSTEMS WITH UNICOS CPC FRAMEWORK

W. Booth, B. Bradu, E. Blanco, M. Quilichini, M. Bes, M. Zimny, R. Barillere,
CERN, Geneva, Switzerland

Abstract

At CERN there are currently 200 ventilation air handling units in production, used in many different applications, including building ventilation, pressurization of safe rooms, smoke extraction, pulsion/extraction of experimental areas (tunnel, cavern, etc.), and the ventilation of the computing centre. The PLC applications which operate these installations are currently being revamped to a new framework (UNICOS CPC [1]). This work began 3 years ago, and we are now in a position to standardize the development of these HVAC applications, in order to reduce the cost of initial development (including specification and coding), testing, and long-term maintenance of the code. In this paper we will discuss the various improvements to the process, and show examples, which can thus help the community develop HVAC applications. Improvements include templates for the "Functional Analysis" specification document, standardized HVAC devices and templates for the PLC control logic, and automatically generated test documentation, to help during the Factory Acceptance Test (FAT) and Site Acceptance Test (SAT) processes.

OVERVIEW

The development of an HVAC (Heating, Ventilation and Air Conditioning) control system follows a standard development cycle including data gathering, requirements specification, application development, implementation, and testing, as outlined below.

DATA GATHERING

The data gathering process collects information on the hardware that is, or will be, on-site. The process control overview is shown in the Process & Instrumentation Diagram (P&ID), see Figure 1. The electrical wiring diagram provides the information on the specific signals available in the PLC. Early on in the migration of cooling and ventilation applications to the UNICOS CPC (Unified Industrial Control System: Continuous Process Control) framework, a standard input/output (IO) list format was adopted, which provides all the pertinent information regarding the IO (e.g. unit name, unit number, actuator, position, hardware type, address, etc), based on the electrical wiring diagram, in a convenient form (Excel spreadsheet), see Figure 2. These two documents form the initial inputs for the control system design and requirements specification.

.

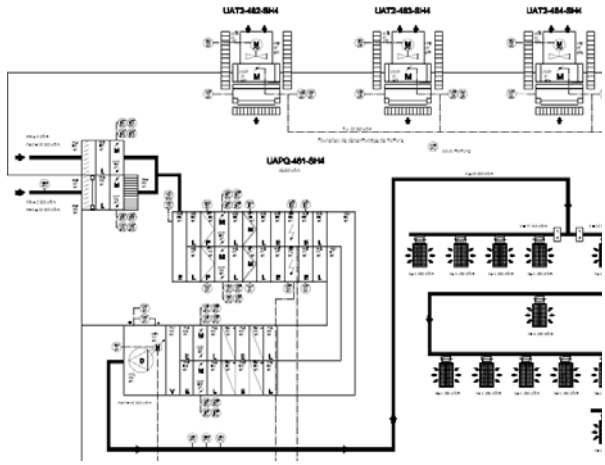


Figure 1: Example Process and Instrumentation Diagram (P&ID)

DeviceIdentification	Unit Name	UAPQ
	Unit Number	481_1
	Actuator	UMRM
	Position	B04
	Generic Name	UBLP
	Special Name	
DeviceDocumentation	Description	DAMPER POSITION EXTERNAL AIR 1 UVUM_B04_481.1
	GMAO	
	Remarks	Attention [2 10] V
FEDeviceIOConfig	Hard. Type	[0 10] V
	Address	PIW128
FEDeviceParameters	Range Min	0
	Range Max	100
	Unit	%

Figure 2: Example IO list format

REQUIREMENTS SPECIFICATION

The requirements of the control system are specified in the functional analysis (FA) document, which defines the functional breakdown of the control system into a hierarchy of individual units as specified in IEC 61512-1 or ANSI/ISA S88 [2]. After 3 years development of Ventilation plants, many different variations were found. Therefore, an analysis of the various ventilation control systems revealed that they are based on two key features: the Stepper (i.e. the state diagram) governing the operation of the individual unit, see Table 1, and the type of temperature regulation, see Table 2.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

APPLICATION OF PyCDB FOR K-500 BEAM TRANSFER LINE*

Pavel Cheblakov#, Sergey Karnaev,
Oidin Khudayberdieva, BINP SB RAS, Novosibirsk, Russia

Abstract

The new injection complex for VEPP-4 and VEPP-2000 e-p colliders is under construction at Budker Institute, Novosibirsk, Russia. The double-direction bipolar transfer line K-500 of 130 and 220 meters length respectively will provide the beam transportation from the injection complex to the colliders with a frequency of 1 Hz. The designed number of particles in the transferred beam is 2*10^10 of electrons or positrons, the energy is 500 MeV. K-500 has dozens of types of magnets, power supplies and electronic devices. It is rather complicated task to store and manage information about such a number of types and instances of entities, especially to handle relations between them. This knowledge is critical for configuration of all aspects of control system. Therefore we have chosen PyCDB to handle this information and automate configuration data extraction for different purposes starting with reports and diagrams and ending with high-level applications and EPICS IOCs' configuration. This paper considers concepts of this approach and shows the PyCDB database structure designed for K-500 transfer line. An automatic configuration of IOCs is described as integration with EPICS.

INTRODUCTION

Every large scientific facility is a complicated construction with a huge amount of hardware, electronic equipment and software. Hardware is controlled by built-in electronics in some cases, it requiring dedicated or universal standalone controllers in other cases. Anyway it results in using of a wide spectrum of electronics of different types from different vendors. Moreover it can rapidly increase the amount of physical protocols types (field busses), interconnections, network equipment, CPU and OS platforms and etc. Consequently, software should provide support for diversity of electronics and thereby hide its complication from user.

The largest facilities (e.g. ITER [1]) include hundreds of local systems and dozens of subsystems. These local systems consist of a set of controllers, interfacing the actuators and sensors, and connected together via network switches to the plant operation network. The total amount of computers can reach one thousand, the total number of signals (wires) and the total number of process variables is estimated at 100.000 and 1.000.000 accordingly.

* This work has been supported by Russian Science Foundation (project N 14-50-00080).
P.B.Cheblakov@inp.nsk.su

PROBLEMS

To maintain so large and diverse system it is necessary to collect and store all relevant information: not only lists of systems, subsystems, network equipment, controllers, instrumentation and control devices but relationships between them as well. We need to know which controller is responsible for corresponding power supply or specialized gauge. We need to know how that controller is connected to corresponding computer and which switch is used. Also we need to know what types of cables are used for interconnection.

Strictly speaking we can continue this list further and further adding more and more information about our facility. This information is vital for broad range of specialists: from storekeepers, designers and electricians to software developers, operators and physicists. Having this data we can create a list of equipment of particular type for record keeping purposes. Or we can simply generate actual cable schedule. Furthermore we can get configuration data for software involved in control.

This information is mostly distributed along many spreadsheets and databases which are weakly connected. This fact is likely to cause data duplication and denormalization so we can get data inconsistency induced by human factor. As a result it is very difficult to change and analyse distributed data both for human and computer.

There are different approaches to store and handle this information. First of all we divide the information into two parts: the first one is used for software configuration and the second one includes everything else. Handling the former part of data is more formalized and there are some practices to operate with it. For example the paper [2] distinguishes centralized and decentralized approaches covering almost all cases and proposes hybrid technique. Even though it turns to a laborious work for every facility [3, 4, 5].

As regards the latter part of the data it is so unstructured and non-formalized that it becomes a great challenge to build up a data-storing system from scratch or even deploy any specialized software system [6].

KNOWLEDGE BASE

To deal with heterogeneity and complexity of data we propose to use a centralized storage which is able to store and handle with ease not only data and its properties but relationships as well. Storing and retrieving relationships is a key aspect for successful functioning of many systems: from equipment procurement and registering to cable scheduling, from GIS-data to software configuration

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

MANAGING A REAL-TIME EMBEDDED LINUX PLATFORM WITH BUILDROOT

J. Diamond#, K. Martin, FNAL*, Batavia, IL 60510, U.S.A.

Abstract

Developers of real-time embedded software often need to build the operating system, kernel, tools and supporting applications from source to work with the differences in their hardware configuration. The first attempts to introduce Linux-based real-time embedded systems into the Fermilab accelerator controls system used this approach but it was found to be time-consuming, difficult to maintain and difficult to adapt to different hardware configurations. Buildroot is an open source build system with a menu-driven configuration tool (similar to the Linux kernel build system) that automates this process. A customized Buildroot [1] system has been developed for use in the Fermilab accelerator controls system that includes several hardware configuration profiles (including Intel, ARM and PowerPC) and packages for Fermilab support software. A bootable image file is produced containing the Linux kernel, shell and supporting software suite that varies from 3 to 20 megabytes large – ideal for network booting. The result is a platform that is easier to maintain and deploy in diverse hardware configurations.

INTRODUCTION

Scientific Linux [2] is a Linux distribution produced by a collaborate effort between Fermilab, CERN and other members of the global Linux community. Scientific Linux has been used successfully as a target platform for data acquisition in the Fermilab accelerator control system but only on commodity PC hardware and not with real-time extensions [3]. When attempting to adopt Scientific Linux to a real-time embedded application some road-blocks were encountered. The first was that Linux is not designed to be a hard real-time operating system. There are however, several options for making Linux real-time and all of these options require a custom kernel be built. At Fermilab, one route chosen was to use the Real-Time Application Interface (RTAI) [4]. The next road-block was the minimum foot-print of the installed system. The minimum Scientific Linux install profile requires 1.5GB of storage. This is a concern because it is desirable to network boot the targets. The requirement for a network boot capability drastically reduces the acceptable footprint size as the entire boot image needs to be transferred over the network. Finally, Scientific Linux is a binary distribution targeted for x86 architecture, thus it would not be possible to use as a platform for ARM and PowerPC targets. These concerns drove the decision to use a custom Linux platform built from the ground-up.

* Operated by Fermi Research Alliance, LLC under Contract No. De-AC02-07CH11359 with the United States Department of Energy

jdiamond@fnal.gov

LINUX "FROM SCRATCH"

The Linux from Scratch website [5] was chosen as the model for the attempt to build a custom Linux platform from source. Linux from Scratch is an on-line resource that explains in step-by-step fashion how to build an entire Linux operating system from source. First the user partitions the target's disk on the build system. Then the necessary source packages (about three dozen) and patches (another dozen) are downloaded onto the build system and unpacked into a special "chroot environment". Then the toolchain packages are built and installed – not a trivial task in its own right. The toolchain is then used to build a complete Linux system from inside of the chroot environment - kernel, boot-loader, shell, etc. Once the system is in place the boot-scripts are installed and configured. Finally, the boot-loader is installed on the target CF disk and an attempt is made to boot the target (things rarely work on the first try).

As illustrated above, the process is exhausting and can take several days to complete. It should be noted that the target platform was x86 so a cross-compile toolchain was not attempted (ARM targets came later). Once developed the root filesystem could be used to boot any target with a similar hardware architecture.

The root filesystems produced with this method were in the range of two to three hundred megabytes. One significant drawback is that a boot image of this size is still too large to network boot. The decision was made to store the kernel, boot-scripts and network configuration on a Compact Flash (CF) disk installed with each target node and the root filesystem on a common network filesystem that could be mounted by each target at boot time. The network filesystem also contains the application programs, libraries and data areas.

One issue that was noticed early on was the failure rate of the CF disks. CF disks are susceptible to failure after many write-cycles. As noted above, application data is written to a network filesystem instead of the CF disk. To reduce the CF write cycles the /tmp and /log partitions were moved to a RAM disk and the system log was redirected to a syslog server. These efforts were able to mitigate the failure rate of the CF disks.

Another significant challenge with this platform was maintainability. Installing new software packages in the root filesystem was easy but deploying to targets already in the field meant physically replacing the CF disks. Furthermore, the root filesystem installed on the network filesystem would creep away from its original state over time without version control system in place.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

LOCAL MONITORING AND CONTROL SYSTEM FOR THE SKA TELESCOPE MANAGER: A KNOWLEDGE-BASED SYSTEM APPROACH FOR ISSUES IDENTIFICATION WITHIN A LOGGING SERVICE*

M. Di Carlo, INAF Osservatorio Astronomico di Teramo, Italy
M. Dolci, INAF Osservatorio Astronomico di Teramo, Italy
G. M. Le Roux, SKA South Africa, Cape Town
R. Smareglia, INAF Osservatorio Astronomico di Trieste, Italy
P. S. Swart, SKA South Africa, Cape Town

Abstract

The SKA Telescope Manager (SKA.TM) is a distributed software application aimed to control the operation of thousands of radio telescopes, antennas and auxiliary systems (e.g. infrastructures, signal processors, ...) which will compose the Square Kilometre Array, the world's largest radio astronomy facility currently under development. SKA.TM, as an "element" of the SKA, is composed in turn by a set of sub-elements whose tight coordination is ensured by a specific sub-element called "Local Monitoring and Control" (TM.LMC).

TM.LMC is mainly focussed on the life cycle management of TM, the acquisition of every network-related information useful to understand how TM is performing and the logging library for both online and offline sub-elements. Given the high complexity of the system, identifying the origin of an issue, as soon as a problem occurs, appears to be a hard task. To allow a prompt diagnostics analysis by engineers, operators and software developers, a Knowledge-Based System (KBS) approach is proposed and described for the logging service.

INTRODUCTION

A log message is the simplest possible storage abstraction which says what happened and when. It is an append-only, totally-ordered sequence of records timestamped. So, a log is not all that different from a file or a table. A file is an array of bytes, a table is an array of records, and a log is really just a kind of table or file where the records are sorted by time.

Since it is a very simple concept, developers tend to underestimate the logging system but logs record *what happened and when* and, for distributed data systems, this can be the only way to find out the origin of an error.

Usually log files are written in a natural language (human readable) and, even if it is very common, this is not the best way to store informations: it does not allow to reason programmatically** about those information.

Building a logging system with a declarative language (for instance prolog) can give the possibility to reason about facts of the communications and operations with an inference engine. For the specific case this document is aimed to, with the adoption of a knowledge-based system approach, TM.LMC could give to SKA engineers, operators and software developers the possibility to ask high level questions to the system in order to understand how a failure came up or simply understand how the system is working.

Information to Log

In order to understand which are the informations to log it is important to make some consideration. In a distributed environment:

- the entities which make up an application are active (processes or agents);
- the interactions are based on message exchange mechanism;
- the process life time is connected to the application life time; the life time of an agent are usually independent from life time of a specific application;
- the logical architecture can be set by different patterns: client-server, peer to peer, etc.
- the middleware realize the physical and logical connection between entities (subsystem, service, object, component, process, agent, etc.).

From an high level point of view the kind of applications like SKA TM define a logical network of interactions which it is composed by different nodes that interact each other and some of them act as coordinator or controller (at least for the online part of the system). So it is important to log:

- *Node† identification,*
- *Node signal‡ declaration,*
- *Node interactions,*
- *Actions and loops.*

* Work supported by the Italian Ministry of University and Research (MIUR)

† In a network, a node is a connection point, either a redistribution point or an end point for data transmissions. In general, a node has programmed or engineered capability to recognize and process or forward transmissions to other nodes.

‡ An information usable by a node in order to control the behaviour of another one or its behaviour in function with the one from another node.

** Doing something programmatically means that you can do it using source code, rather than via direct user interaction or a macro

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

DRAMA 2 - AN EVOLUTIONARY LEAP FOR THE DRAMA ENVIRONMENT FOR INSTRUMENTATION SOFTWARE DEVELOPMENT

T. Farrell, K. Shortridge, Australian Astronomical Observatory, Australia

Abstract

The DRAMA Environment provides an API for distributed instrument software development. It originated at the Anglo-Australian Observatory (now Australian Astronomical Observatory) in the early 1990s, in response to the need for a software environment for large distributed and heterogeneous systems, with some components requiring real-time performance. It was first used for the AAOs 2dF fibre positioner project for the Anglo-Australian Telescope. DRAMA is used for most AAO systems and is or has been used at various other observatories looking for a similar solution. Whilst DRAMA has evolved and many features were added, the overall design has not changed. It was still a largely C language based system, with some C++ wrappers. It did not provide good support for threading or exceptions. Ideas for proper thread support within DRAMA have been in development for some years, but C++11 has provided many features that allow a high quality implementation. We have taken the opportunity provided by C++11 to make significant changes to the DRAMA API, producing a modern and more reliable interface to DRAMA, known as DRAMA2.

INTRODUCTION

The DRAMA API [1] remains the AAO's primary tool for constructing complex instrumentation systems and has been/is being used by various other observatories. With an approach based on the older Starlink ADAM Environment [2], it implements a tasking model; with each named task responding to named messages of a number of different types. In a DRAMA "System", tasks can run across different hosts in a heterogeneous environment. DRAMA was implemented from about 1992 and was designed to be highly portable at a time before ANSI C was available on all machines of interest. It has been run on many flavours of UNIX/Linux, VMS, VxWorks and MS Windows, and provided the ability to write soft[1] real-time applications and with good performance on, for example, 30Mhz 68020 CPUs. The flexibility allowed systems as complex as the AAO's 2dF system [3] to be implemented, making use of the most appropriate hardware for each job across a distributed system.

Most work is a DRAMA task is done in response to "Obey" messages – in effect, command messages; implementing "Actions". The design approach implements co-operative multi-tasking; multiple actions can be running at the same time but must deliberately return control to the DRAMA message reading loop between events to allow other actions to run and for the action itself to be "Kicked" – sent a message to change its behaviour in some way (typically, but not always, to cancel the action cleanly). The approach has worked well and a strongly objected-oriented task design approach was implementable for tasks written in C.

Attempts were made starting about 1994 to implement C++ interfaces for DRAMA, but the results were relatively poor and various different approaches were tried. One of the early issues was the poor portability of early C++ compilers, some features such as templates and exceptions were not reliably implemented and were not portable. Another was that we were still learning the best approaches to use.

Whilst DRAMA tasks using threads of various types have been implemented over the years, DRAMA itself has not supported using threads, with its own co-operative multi-tasking technique sufficient in most cases being more portable then threads were. In the C API, task authors must work around DRAMA when using threads; but in recent times, many libraries for component control have presumed threads are available and thread support has become widespread and is presumed to be available by most software engineers. We had been working on designs for proper thread support for DRAMA over some years, but had not yet implemented it.

C++11 [4] was a major revamp to the C++ language: Threads are now supported using a well thought out approach, by the compilers and standard libraries; Many new features are provided by C++11 that assist library implementers to construct quality interfaces; compilers of interest (GCC and Clang in particular) have implemented the full feature set on machines of interest (Linux and Mac OS X). We have taken advantage of the upgrade of C++ to implement DRAMA2, which will simplify writing and maintaining complex DRAMA tasks.

BASIC DRAMA

To understand DRAMA2, a quick introduction to DRAMA is needed. A DRAMA task executes a message receive loop that dispatches control to event/message handlers when messages arrive. A "Path" is a connection to another task, which is opened as required by application specific code and then used to send messages. There are various types of messages sent between tasks, which may be running across various machines on the network.

1 "Soft" Real-Time: Don't bet a life on an interrupt response, but good enough for ground based astronomy.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

A MODULAR SOFTWARE ARCHITECTURE FOR APPLICATIONS THAT SUPPORT ACCELERATOR COMMISSIONING AT MedAustron

M. Hager*, M. Regodic#, EBG MedAustron, Wiener Neustadt, Austria

Abstract

The commissioning and operation of an accelerator requires a large set of supportive applications. Especially in the early stages, these tools have to work with unfinished and changing systems. To allow the implementation of applications that are dynamic enough for this environment, a dedicated software architecture, the Operational Application (OpApp) architecture, has been developed at MedAustron. The main ideas of the architecture are a separation of functionality into reusable execution modules and a flexible and intuitive composition of the modules into bigger modules and applications. Execution modules are implemented for the acquisition of beam measurements, the generation of cycle dependent data, the access to a database and other tasks. On this basis, Operational Applications for a wide variety of use cases can be created, from small helper tools to interactive beam commissioning applications with graphical user interfaces. This contribution outlines the OpApp architecture and the implementation of the most frequently used applications.

INTRODUCTION

The heart of MedAustron is a synchrotron-based accelerator. The accelerator provides the possibility to generate a large range of different ion beams, for example Proton and Carbon beam with different beam sizes and hundreds of different energies. A backside of synchrotron-based accelerators is the high number of components and the complex control of the beam. The commissioning of such an accelerator is a laborious task that requires support by software to be executed efficiently.

One challenge in the development of such applications is the agility of the commissioning process. While the commissioning progresses, new components are integrated, existing ones change and the understanding of the accelerator and the beam behavior deepens. Software that supports the process has to be easily adaptable to the evolving environment. Previous architectural concepts turned out to be too restrictive for these demands. Therefore, a new, modular architecture has been developed at MedAustron, the OpApp architecture.

APPROACH

The OpApp framework is connected to the MedAustron Accelerator Control System (MACS), [1]. With the connection to MACS, OpApps can: A) Send commands to the systems and devices of the accelerator B) Apply device settings and retrieve state information. C) Request beam cycles D) Receive measurements and timing information.

The framework is also connected to a database. OpApps use the database to store data related to the beam generation as well as acquired measurements and accelerator configuration. OpApps also retrieve and analyze stored data.

Functionality

Main domain of Operational Applications is the beam commissioning. OpApps can compute settings, like currents and voltages, for all accelerator devices based on the optical setup of the accelerator and the desired beam characteristics. OpApps can also apply the settings, request beam cycles and measure the characteristics of the generated beam. With this, OpApps can automatize many aspects of the measurement-based beam commissioning workflows. The role of Operational Applications in the workflow is depicted in Figure 1.

OpApps contribute to the processing and analysis of measurements but usually leave complex analysis tasks to dedicated tools. To provide data to other tools, OpApps can generate files in a variety of different formats.

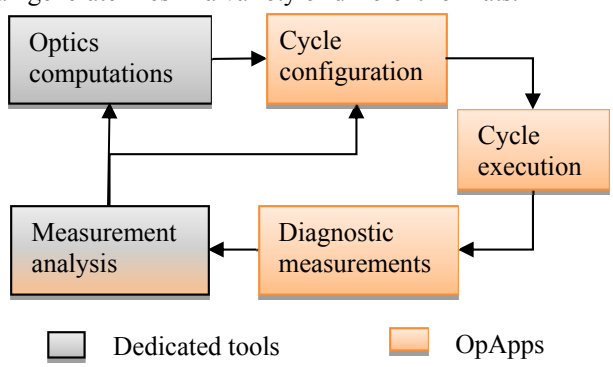


Figure 1: The measurement-based beam commissioning workflow

Although the OpApp concept was developed with beam commissioning in mind, OpApps are by far not limited to it. Operational Applications can also be used for:

- Quality Assurance (QA) - OpApps that acquire measurements and compare them with stored reference data can be used for a regular QA of the beam characteristics
- Configuration Management - With a set of database related execution modules OpApps can, for example, help the user to import device specifications into the database or export data for the Control System
- Accelerator and Beam Monitoring - OpApps can be used to acquire accelerator and beam data in the

* markus.hager@medaustron.at
milovan.regodic@medaustron.at

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

SOLVING THE SYNCHRONIZATION PROBLEM IN MULTI-CORE EMBEDDED REAL-TIME SYSTEMS

F. Hoguin, S. Deghaye, CERN, Geneva, Switzerland

Abstract

Multi-core CPUs have become the standard in embedded real-time systems. In such systems, where several tasks run simultaneously, developers can no longer rely on high priority tasks blocking low priority tasks. In typical control systems, low priority tasks are dedicated to receiving settings from the control room, and high priority real-time tasks, triggered by external events, control the underlying hardware based on these settings. Settings' correctness is of paramount importance and they must be modified atomically from a real-time task point of view. This is not feasible in multi-core environments using classic double-buffer approaches, mainly because real-time tasks can overlap, preventing buffer swaps. Other common synchronization solutions involving locking critical sections introduce unpredictable jitter on real-time tasks, which is not acceptable in CERN's control system. We present a lock-free, wait-free solution to this problem based on a triple buffer, guaranteeing atomicity no matter the number of concurrent tasks. The only drawback is potential synchronization delay on contention. This solution has been implemented and tested in CERN's real-time C++ framework.

FROM SINGLE-THREADED TO MULTI-THREADED EMBEDDED SYSTEMS

In typical control systems, accelerators settings are modified either manually by operators using graphical user interfaces or by high-level systems computing hardware settings from high-level values. These settings are then sent by the high-level applications to the computers in charge of the hardware real-time control. This control is done by real-time tasks in a limited time frame following an external trigger. A real-time task is nothing more than a piece of code executed with real-time priority. These tasks typically perform computations based on settings and drive the hardware with the values obtained. Hardware settings can be interdependent and modifications must be applied in a single operation. The real-time tasks must work with consistent sets of settings; from their point of view, settings must be modified atomically.

In single core embedded systems, developers can rely on the determinism of a real-time scheduler to guarantee consistency of settings with a simple double buffer system; one for active values (accessed by real-time tasks) and another one for pending values i.e. just modified values not yet accessible to real-time tasks. When the modification is done, the set of pending values is consistent and buffers may be swapped. A low priority task swaps pointers to the buffers atomically; the pending buffer becomes active, and vice versa. This is guaranteed to be safe provided no real-time task is ever in a waiting state in the middle of its execution, ensuring that the buffer swapping task can never be executed while a real-time task is being executed.

As multi-core embedded systems become the norm, such assumptions cannot be made anymore. At CERN, most of the control system embedded systems use 2-core processors. In such setups, the buffer swapping task can be executed concurrently to a real-time task. As a result, swapping pending and active buffers without further checks could lead to a real-time task reading inconsistent settings. An example is shown in table 1.

Table 1: Example of Naive Buffer Swap with Multi-core CPU

Real-time task	Buffer swapping thread	Buffer 1	Buffer 2
Reads voltage => 200V	Swap triggered…	V = 200, A = 100	V = 10, A = 1000
Busy computing…	Swaps buffers	V = 200, A = 100	V = 10, A = 1000
Reads current => 1000A		V = 200, A = 100	V = 10, A = 1000
Tells power supply: 200V, 1000A		V = 200, A = 100	V = 10, A = 1000

EXPLORED AND ABANDONED SOLUTIONS

In order to solve this problem, several solutions were explored. The goal was to implement an algorithm that fulfils the following requirements:

1. The solution shall be real-time compliant (in particular, no dynamic memory allocation is allowed).
2. Real-time tasks shall read consistent setting values throughout their execution.
3. Time interval between a trigger and the execution of the corresponding real-time task shall be constant and below 5 milliseconds (a jitter of 10% is acceptable).
4. Pending setting values shall be made available to real-time tasks within a reasonable time frame.

Snapshot of the Setting Values

A possible solution is to make a snapshot of the setting values just as a real-time task is about to start. This snapshot is private to that execution of the task. While this would work, it is not real-time compliant as memory would be dynamically allocated to copy the setting values. Even a fixed size memory pool could be exhausted given enough simultaneous real-time tasks. Also, the jitter between the trigger and the task execution

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

EPICS V4 EVALUATION FOR SNS NEUTRON DATA*

K.U. Kasemir, G.S. Guyotte, M.R.Pearson, ORNL, Oak Ridge, TN37831, USA

Abstract

Version 4 of the Experimental Physics and Industrial Control System (EPICS [1]) toolkit allows defining application-specific structured data types (pvData) and offers a network protocol for their efficient exchange (pvAccess). We evaluated V4 for the transport of neutron events from the detectors of the Spallation Neutron Source (SNS) to data acquisition and experiment monitoring systems. This includes the comparison of possible data structures, performance tests, and experience using V4 in production on a beam line.

MOTIVATION

On SNS beam lines, each neutron event consists of a pixel ID that identifies the location of the detector where a neutron was observed, and a time-of-flight measurement that describes when the neutron was detected relative to the most recent beam pulse. Depending on the beam line and its specific configuration, event rates can reach a few million events per seconds.

This neutron event information needs to be transferred from detectors to processing stages that provide users of the experiment with visual feedback, accumulate information that allows for the automation of the experiment, and finally stream the events into a data collection pipeline for long-term storage of the experiment data.

The original SNS beam line data acquisition software used a locally developed UDP/IP-based network protocol to transmit neutron event information [2]. The limited performance and reliability of this protocol necessitated an update of the overall data acquisition software [3].

EPICS V4

Based on a proven track record for control of the SNS accelerator, EPICS had been chosen as a toolkit for the beam line control system upgrade. While Channel Access [4], the original EPICS network protocol, has a well defined and functional set of data objects, this set is fixed to types suitable for describing a single data point, for example a temperature reading or voltage set point.

*This manuscript has been authored by UT-Battelle, LLC under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan(http://energy.gov/downloads/doe-public-access-plan).

EPICS V4 [5] is an addition to the EPICS toolkit that introduces an alternative to the existing EPICS V3 data types and network protocol.

pvData

pvData is the EPICS V4 library for structured data. It can describe such data in an operating-system independent way, hold it in memory, and copy complete or partial data containers. The data can include time stamps, numeric values, enumerated data, text and alarm information. Data can also be assembled into arrays and structures.

Normative Types

pvData allows clients to define nearly arbitrary data structures, which is ideal for packaging site-specific information. At the same time it limits the interoperability of applications. *Normative Types* are a set of agreed-upon pvData types that all implementers of V4 applications are encouraged to support. All original EPICS V3 data types are described as Normative Types, allowing for an eventual transition from V3 to V4. In addition, data types like N-dimensional images or statistical samples that are often used in higher-level control system applications are available as Normative Types.

pvAccess

pvAccess is the V4 network protocol that allows for the exchange of pvData. It is conceptually similar to V3 Channel Access, using UDP/IP for channel name resolution, and then establishing a TCP/IP connection between pvAccess servers and clients to exchange data. It supports basic read and write access. A subscription mode efficiently updates clients on changes in the data by only transferring the modified structure elements. Finally, a combined write/read mode supports remote service calls by atomically sending parameters, awaiting the execution of the remote service, then returning the result.

Both pvData and pvAccess have been implemented in C++ and Java, with additional bindings for Python [6].

SNS NEUTRON DATA

SNS neutron data consists of a list of pixels and time-of-flights as already described, combined with a sequential pulse number and the proton charge of the accelerator pulse that generated these neutrons.

The following pvData structure would be a direct representation:

```
// Sequential pulse number
uint64   pulse
```

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

CCLIBS: THE CERN POWER CONVERTER CONTROL LIBRARIES

Q. King, K. Lebioda, M. Magrans De Abril, M. Martino, R. Murillo, A. Nicoletti,
CERN, Geneva, Switzerland

Abstract

Accurate control of power converters is a vital activity in large physics projects. Several different control scenarios may coexist, including regulation of a circuit's voltage, current, or field strength within a magnet. Depending on the type of facility, a circuit's reference value may be changed asynchronously or synchronously with other circuits. Synchronous changes may be on demand or under the control of a cyclic timing system. In other cases, the reference may be calculated in real-time by an outer regulation loop of some other quantity, such as the tune of the beam in a synchrotron. The power stage may be unipolar or bipolar in voltage and current. If it is unipolar in current, it may be used with a polarity switch. Depending on the design, the power stage may be controlled by a firing angle or PWM duty-cycle reference, or a voltage or current reference. All these cases are supported by the CERN Converter Control Libraries (CCLIBS). These open-source C libraries include advanced reference generation and regulation algorithms. This paper introduces the libraries and reviews their origins, current status and future.

INTRODUCTION

The CERN converter control libraries are a collection of libraries written in C, available under the GNU Lesser General Public License from https://github.com/cclibs.

The libraries were started in 2010 by extracting the measurement calibration, function generation and current regulation code from the operational software used to control the power converters in the CERN LHC. This software was originally deployed in the LHC magnet test facility in 2000 [1,2], so the libraries built on ten years of experience with the high-precision control of current in magnet circuits [3].

Figure 1 shows an overview of the architecture of a power converter controller based on CCLIBS. The long-term objective is to move all the re-usable generic software components into CCLIBS and to leave the application with just the hardware specific components.

In 2012, the new libraries [4] were deployed into operation on two different hardware platforms [5] and a third platform was added the following year.

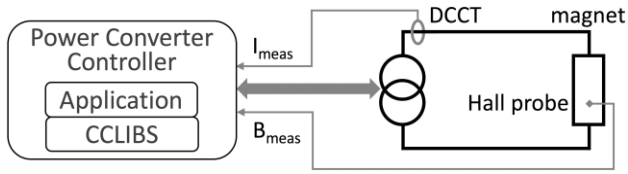


Figure 1: The role of CCLIBS in converter control

In 2013, a second version of the libraries was started with the objective to support regulation of magnetic field measured by a hall probe. Hall probes have more noise than current transducers (DCCTs), so more advanced anti-aliasing filtering is needed. This in turn requires more advanced algorithms for the computation of the coefficients for the RST regulator, to accommodate the additional delay from the filters.

Then in 2015, a new reference manager library, libref, was created to support the three different operating modes used with magnet circuits at CERN. This is well advanced and will be deployed into operation in 2016, along with two more libraries, libsig and liblog. Several other features will also be added, including voltage regulation for thyristor converters.

THE LIBRARIES

Taken together, the Converter Control Libraries have too many features to cover them all in this short paper. Furthermore, the first versions of libfg and libreg were described in [4] so this paper will report some of the improvements made to libfg and libreg as well as the main features of the new libraries and test programs.

Libcal: Calibration

The calibration library was part of CCLIBS from the start. It provides a three-point calibration scheme for ADCs and DCCTs, with second order temperature compensation. This accommodates an offset and separate positive and negative gain errors. The library also supports three-point DAC calibration and a simple scaling for hall probes and voltage probes (i.e. no offset and just one gain that applies to both positive and negative measurement values). The library will be optimised for better performance in 2016.

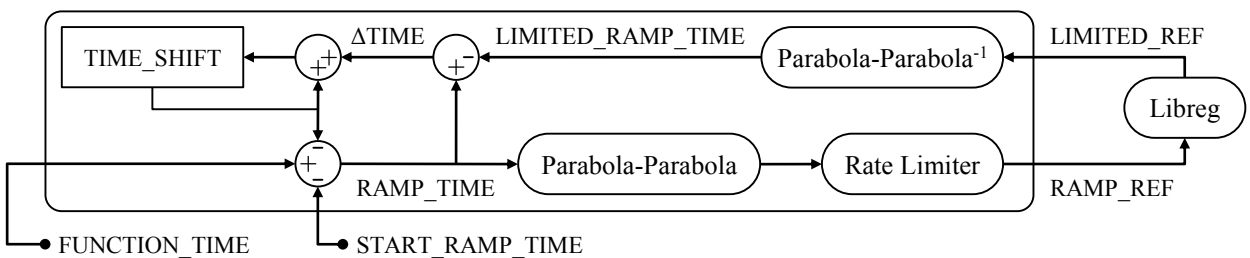


Figure 2: The RAMP function algorithm

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

MULTI-HOST MESSAGE ROUTING IN MADOCA II

T. Matsumoto#, Y. Furukawa, K. Okada, JASRI/SPring-8, Hyogo, 679-5198, Japan

Abstract

MADOCA II is the next generation of the Message and Database-oriented Control Architecture (MADOCA) that was implemented in the control systems of the SPring-8 and the SPring-8 Angstrom Compact Free Electron Laser (SACLA) data acquisition (DAQ) systems in 2013 [1]. In 2014, SACLA was equipped with a second beamline (BL2) to increase the capacity of experiments. Messaging over a firewall network is required for beamline control, and an ad hoc socket application was used for this because MADOCA II does not have the requisite functionality. However, this caused problems and increased the complexity of the control system. In this paper, we propose multi-host message routing in MADOCA II to get rid of the socket application In SACLA DAQ systems, multi-host message routing is used for sophisticated access control. We introduce a master server that mediates control messages between operator workstations and equipment management servers. Since access control can be centralized to the master server, reliable operation was achieved by avoiding the effects of the accidental modification of DAQ settings by end users. We implemented multi-host message routing in the SACLA DAQ system in September 2014, and it has been stably operating since then.

INTRODUCTION

As the next generation of the Message and Database-oriented Control Architecture (MADOCA), MADOCA II was successfully implemented to control systems for SPring-8 and the SPring-8 Angstrom Compact Free Electron Laser (SACLA) data acquisition (DAQ) system in 2013, as reported at the last meeting of the International Conference on Accelerator and Large Experimental Physics Control Systems (ICALEPCS) [1]. Like its predecessor, MADOCA II is based on a message-oriented control scheme. However, messaging in this architecture has been made more flexible by replacing the messaging protocol with the ZeroMQ socket library [2] and reconfiguring the messaging scheme. In MADOCA II, data of varying length, such as image data can be attached to messages, and control applications can be used in the Windows operating system environment. Such features were utilized in several control applications in SPring-8 [3] [4].

However, MADOCA II needed to be extended when a second beamline (BL2) was equipped to the SACLA system in 2014. There are several network layers in the control system for experimental users, data acquisition, and the beamline equipment. We sometimes need control over the firewall network, which requires messaging via an intermediate host, as shown in Figure 1.

MADOCA II is designed to facilitate communication between hosts. To implement message routing among more than two hosts, we used an ad hoc socket application to mediate messages. However, we found that the application stalled at times. This was because the application processes messages sequentially, and is affected when a messaging procedure takes a while. Moreover, the application renders the control procedure complex, making it difficult to understand when and why control troubles occurred. To solve these problems, we developed multi-host message routing in MADOCA II.

In the following section, first, we provide an account of the implementation of multi-host routing in MADOCA II and its application to the SACLA DAQ system. We conclude with a summary of our findings.

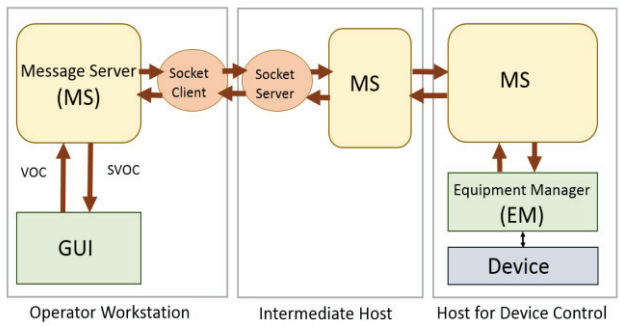


Figure 1: Messaging via an intermediate host in MADOCA II. An ad hoc socket application is used to mediate a message.

IMPLEMENTATION OF MULTI-HOST MESSAGE ROUTING IN MADOCA II

In this section, we describe the implementation of multi-host message routing in MADOCA II. We first explain the messaging procedure between two hosts, followed by an account of how the scheme is extended to multi-host message routing. Some technical issues, such as access control, prevention of messaging loops, and RTT (round-trip time) during messaging, are also addressed.

Message Routing Between Hosts

Figure 2 shows an example of messaging between hosts. A message command is exchanged between a graphical user interface (GUI) on host-b and an equipment manager (EM) on host-a through a Message Server (MS) in each host. In MADOCA II, messages are based on text and composed of a character string in subject/verb/object/complement (S/V/O/C) syntax. In case of messages from a GUI, “put” (V) represents the

#matumot@spring8.or.jp

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

FLOP: CUSTOMIZING YOCTO PROJECT FOR MVMExxxx POWERPC AND BEAGLEBONE ARM

L. Pivetta, A.I. Bogani, R. Passuello
Elettra Sincrotrone Trieste, Trieste, Italy

Abstract

During the last fifteen years several PowerPC-based VME single board computers, belonging to the MVMExxxx family, have been used for the control system front-end computers at Elettra Sincrotrone Trieste. Moreover, a low cost embedded board has been recently adopted to fulfill the control requirements of distributed instrumentation. These facts lead to the necessity of managing several releases of the operating system, kernel and libraries, and finally to the decision of adopting a comprehensive unified approach based on a common codebase: the Yocto Project. Based on Yocto Project, a control system oriented GNU/Linux distribution called Flop has been created. The complete management of the software chain, the ease of upgrading or downgrading complete systems, the centralized management and the platform-independent deployment of the user software are the main features of Flop.

INTRODUCTION

Since the last decades, technology evolution and component obsolescence are driving a renewal that involves also particle accelerator control system platforms. Even industrial grade products well-known for the long lifetime support, such as VME, suffer these issues. Moreover, the evolution of application software, especially when dealing with new compiler features or when requiring new system library functionalities, often involves an update of the operating system software. Several models of PowerPC-based VME single board computers are in use at Elettra, running different distributions/releases of the GNU/Linux operating system [1–3]. Furthermore, the BeagleBone, a low cost embedded board based on the ARM microprocessor, has been recently introduced, as well as a low-power high-performance INTEL Bay Trail Celeron based Soc platform to interface a specific instrument via USB.

To address all these platforms in an effective and straightforward way, ensuring the consistency of the operating system, the device drivers and the application software, a comprehensive approach is desirable.

GNU/LINUX

The usual approach to GNU/Linux consists in the installation of one of the many available distributions. Yet, most of them do not support all the required architectures and instruction sets and, very often, rely on a generic instruction set to achieve improved compatibility at the expense of the performance.

Using distinct distributions for different architectures, which is the current scenario at Elettra, entails a number of possible issues, among which different releases of the system libraries, of the system initialization, e.g System V init rather than upstart or systemd, and of the filesystem hierarchy. Furthermore, diverse bugfix policy is often in place, making the system and platform adimistration quite onerous in the medium/long term.

Sticking to a specific GNU/Linux distribution release, conversely, exposes to the risk of outdated software, especially concerning the kernel, the system libraries and the compiler. Also, recent hardware is usually poorly supported, or even unsupported, by old distributions without a big effort in back-porting.

THE YOCTO PROJECT

The Yocto Project [4] is an open source collaboration, aimed at providing templates, tools and methods to create custom GNU/Linux based systems for embedded hardware, regardless of the architecture. Established in 2010, it involves many hardware manufacturers, open-source operating system vendors and electronics companies. The Yocto Project provides a complete embedded Linux development environment, with tools, metadata and documentation for free, including core system component recipes provided by the OpenEmbedded project. Specific platform support is provided by the Board Support Package (BSP), for which a standard format has been developed. In addition to the very effective command line tools, the Yocto Project provides an Eclipse IDE plugin and the Hob graphical user interface.

FLOP

Starting from the Yocto Project release 1.8, a control system oriented GNU/Linux distribution named Flop has been created. The current Flop release is 0.5, based on the 3.14 Linux kernel, glibc 2.21, and systemd 219; the compiler in use is gcc release 4.9.2. The TANGO Controls framework is also included, featuring OmniORB 4.1.6, ZeroMQ 3.2.4 and TANGO 8.1.2.

Flop Recipes

In the Yocto Project, the key mechanism that allows to define a specific platform support is referenced as recipe. Each recipe specifies the procedures to select and compile software packages. Recipes can be grouped in layers, and multiple layers can be specified to introduce dependencies. Recipes, usually written as shell scripts o Python scripts, make use of common variables that can be accessed by

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

PHYSICS APPLICATION INFRASTRUCTURE DESIGN FOR FRIB DRIVER LINAC*

G. Shen#, E. Berryman, Z. He, M. Ikegami, D. Liu, D. Maxwell, V. Vuppala,
FRIB, Michigan State University, East Lansing, 48864 USA

Abstract

FRIB (Facility for Rare Isotope Beams) is a new heavy ion accelerator facility to provide intense beams of rare isotopes, and currently under active construction at MSU (Michigan State University). Its driver linac accelerates all stable ions up to uranium, and targets to provide a CW beam with the energy of 200MeV/u and the beam power of 400 kW. A new software infrastructure for high level control and physics application is under development to support the beam commissioning and operation, which is a 3-tier structure consisting of upper level, middle layer, and low level respectively. The detailed design and its current status are described in this paper.

INTRODUCTION

FRIB [1] is a new project under cooperative agreement between US Department of Energy and MSU. It is under construction on the campus of MSU and will be a new national user facility for nuclear science. Its driver accelerator is designed to accelerate all stable ions to energies > 200 MeV/u with beam pow er on the target up to 400 kW as shown in Table 1 [2].

Table 1: FRIB Driver Accelerator Main Parameters

Parameter	Value	Unit
Primary beam ion species	H to ^{238}U	
Beam kinetic energy on target	> 200	MeV/u
Maximum beam power on target	400	kW
Beam duty factor	100	%
Beam current on target (^{238}U)	0.7	emA
Beam spot size on target (90%)	1	mm
Driver linac beam-path length	517	m
Average uncontrolled beam loss	< 1	W/m

The layout of FRIB is as shown in Fig. 1, which consists of two ECR (Electron Cyclotron Resonance) ion sources, a low energy beam transport with a pre-buncher and a chopper, a RFQ (Radio Frequency Quadrupole) linac, LS1 (Linac segment 1) to accelerate the beam up to > 15 MeV/u, LS2 and LS2 to accelerates the beam > 200 MeV/u, two folding segments to confine the footprint and facilitate beam collimation, and a beam delivery system to transport to the target. The beam is stripped to higher charge states after LS1 section.

*Work supported by the U.S. Department of Energy Office of Science under Cooperative Agreement DE-SC0000661
#shen@frib.msu.edu

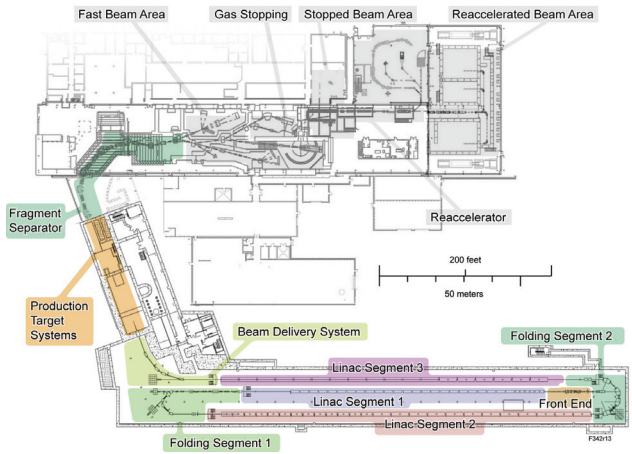


Figure 1: Layout of the FRIB driver accelerator, target and fragment separator as shown in colored area.

ARCHITECTURE

A distributed high level controls environment has been adopted for its beam commissioning and operation as shown in Fig. 2. It is a 3-tier structure consisting of upper level, middle layer, and low level respectively.

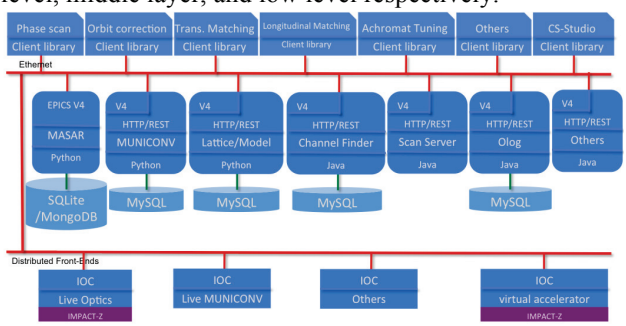


Figure 2: FRIB high level controls system architecture.

Since EPICS (Experimental Physics and Industry Control System) has been adopted as FRIB hardware control framework, the low level is represented as an EPICS IOC (Input Output Controller) for integration purpose. Some services are implemented in the low level layer, the middle layer consists some general services, and high level is mainly for OPI (Operator Interface).

LOW LEVEL SERVICE

At FRIB, the low level service is implemented as an EPICS IOC. One typical service is a virtual accelerator, which provides emulated accelerator behaviour such as setting and reading back a power supply setting, and reading beam orbit at each BPM position.

The architecture of FRIB virtual accelerator is as shown in Fig. 3. It uses IMPACT-Z [3] as underneath simulation engine. During first initialization, it reads

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

LabVIEW EPICS PROGRAM FOR MEASURING BINP HLS OF PAL-XFEL*

Hyojin Choi†, Kwang Won Seo, Youngjin Suh, Kyehwan Gil, Seung Hwan Kim, Heung-Sik Kang,
Department of Accelerator, PAL-XFEL, Pohang, Korea

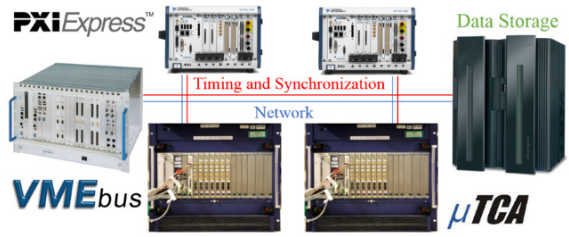
Abstract

In PAL-XFEL, a 4th generation light source, the HLS (Ultrasonic-type Hydrostatic Levelling System) developed at BINP (Budker Institute of Nuclear Physics) in Russia was installed and operated in all parts of PAL-XFEL in order to maintain observations of the vertical change building floor by the ground sinking and uplifting. For this, a HLS measuring program was written using NI LabVIEW and an EPICS IOC Server was built using the CA Lab which has been developed at BESSY (Berlin Electron Storage Ring Society for Synchrotron Radiation) in Germany. The CA Lab was improved and verified in order to confirm that it could support EPICS BASE libraries V3.14.12, and EPICS CA Client and that the EPICS IOC Server could be easily constructed by CA Lab in a 64-bit LabVIEW. This made Multi-core CPU (Multi-core Processor / Multi-thread Program) resource of 64bit Computer System (64bit Hardware PC, 64bit Windows OS, 64bit LabVIEW Multi-thread Programming) to be 100 percent utilized. This study proposes a configuration process for the HLS measuring program algorithm and a building process for the EPICS IOC Server by using CA Lab.

INTRODUCTION

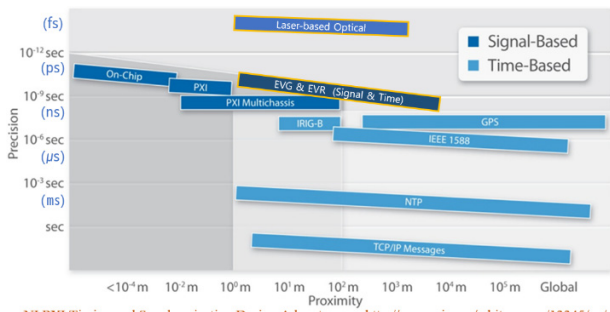
As devices created due to big science and big technology have been placed far away from control systems in terms of space and geography, advanced technology, including that for a distributed control system has become necessary (7.7 Control and Timing [1]). Various dispersed control platforms have been installed for PAL-XFEL like Figure 1. Owing to recent rapid developments in information technology (IT), such as advanced robot technology, Internet of things (IoT), big data and artificial intelligence, it is easy to establish such a new system through buying a distributed control system device developed by diverse companies. Engineers in many accelerator research institutes around the world have been engaged in efforts to advance effective technological developments by sharing technology through the open platforms of control systems and combining their experience and knowledge [2].

As shown in Figure 2, in order to synchronize an ordinary distributed control system, the time-based (time-stamp) synchronization of the entire big scientific device should be effected, first of all. After this, the signal-based synchronization of a module inside each device or the device should be achieved. Regarding RF reference of FEL and pump-probe experimental devices of beamline, a jitter and drift need a synchronization system of the sub-10 rms femtosecond level [3], and a laser-based optical synchronization system for the highest precision of synchronization has been developed to be applied to them (7.8 Femtosecond precision optical synchronization [1])[4].



① Establish a distributed control system.
② Design a high-performance and high-bandwidth network system suitable for a distributed control system and establish it.
③ Install a storage and backup system that can perform high-speed processing of large data in order to store large measurement control data without data loss.
④ Install a timing and synchronization system to get synchronized measurement data.

Figure 1: Construction of a distributed control system.



NI PXI Timing and Synchronization Design Advantages http://www.ni.com/white-paper/13345/en/

Figure 2: Signal-based and time-based synchronization architecture performance.

Timing and synchronization technologies correlate events in time, which is necessary to perform coordinated activities. For software to orchestrate these coordinated activities, the program needs to be synchronized and include a concept of time. As shown in Figure 3, hardware and application programs can be accurately synchronized only when the timing structure is within all systems.

In order to operate big scientific devices stably for a long time, measurement control devices that operate individually also should be highly reliable and available to be operated for a long time without problems occurring. For this purpose, stable and reliable hardware platforms and operating systems (OS) should be selected.

*Work supported by Ministry of the Science, ICT and Future Planning
†choihyo@postech.ac.kr

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

PvaPy: PYTHON API FOR EPICS PV ACCESS*

S. Veseli, Argonne National Laboratory, Argonne, IL 60439, USA

Abstract

As the number of sites deploying and adopting EPICS Version 4 grows, so does the need to support PV Access from multiple languages. Especially important are the widely used scripting languages that tend to reduce both software development time and the learning curve for new users. In this paper we describe PvaPy, a Python API for the EPICS PV Access protocol and its accompanying structured data API. Rather than implementing the protocol itself in Python, PvaPy wraps the existing EPICS Version 4 C++ libraries using the Boost.Python framework. This approach allows us to benefit from the existing code base and functionality, and to significantly reduce the Python API development effort. PvaPy objects are based on Python dictionaries and provide users with the ability to access even the most complex of PV Data structures in a relatively straightforward way. Its interfaces are easy to use, and include support for advanced EPICS Version 4 features such as implementation of client and server Remote Procedure Calls (RPC).

INTRODUCTION

EPICS Version 4 (EPICS4) [1] extends Version 3 [2] with features like support for complex data structures and service oriented architecture, optimized data transfers, as well as support for high level data and image processing. It also comes with a comprehensive set of C++ and Java APIs. However, what has been missing until recently is support for scripting languages. PvaPy aims to fill that gap by providing a Python API for the EPICS PV Access (PVA) protocol.

Rather than providing a direct Python implementation of the PVA protocol, PvaPy wraps EPICS4 C++ code using the Boost.Python [3] framework, a C++ library that enables seamless interoperability between C++ and Python. The main advantage of this approach is that it allows us to build on the existing EPICS4 code base and functionality, which significantly reduces PvaPy development effort.

BUILD PROCESS

Prerequisites for building PvaPy from sources [4] include the following:

- EPICS Base (v3.14.12.x, or v3.15.x) [5]
- EPICS4 C++ release (v4.4.0 or v4.5.0) [6]
- Python development header files/libraries (v2.6.x or v2.7.x) [7]
- Boost (v1.41.0 or later); installation must include the Boost.Python library [3]
- Standard set of GNU development tools (gcc, make, autoconf, etc.) [8]
- Sphinx (Python Documentation Generator) [9]; this is an optional package, generating documentation at build time is not essential.

Except for EPICS Base and the EPICS4 C++ release, all software dependencies listed above are typically included in most Linux operating system (OS) distributions. PvaPy has not been built or tested on Microsoft Windows.

PvaPy utilizes the standard EPICS build infrastructure [10]. However, unlike most EPICS modules, it also offers the possibility of configuring the software build automatically, using the GNU Autoconf [11] and a set of M4 [12] macros. Automated configuration determines compiler flags appropriate for the given operating system, and for the specific versions of Boost and Python that are installed on the build machine. Configuration scripts also determine the PvaPy API version that is suitable for the particular version of EPICS4 release, as well as prepare user environment setup scripts. User setup scripts modify PYTHONPATH environment variable so that PvaPy's "pvaccess" module can be imported within Python scripts or for interactive usage.

SOFTWARE FEATURES

PvaPy provides C++ code which calls the EPICS4 C++ libraries and defines a set of high-level classes for data objects, exceptions, and client/server interfaces. Those classes and their interfaces are exposed to users as the Python "pvaccess" module using the Boost.Python framework. PvaPy also defines a number of low-level utility and helper classes that are either required by EPICS4 APIs, or handle things like conversion between various Python and EPICS4 data structures. Note that the new high-level PVA Client C++ module [13] (available as part of the EPICS4 v4.5.0 release) greatly simplifies EPICS4 client interfaces and significantly reduces the number of internal classes that are implemented in PvaPy.

PvaPy Objects

EPICS4 C++ data types and modelling APIs are part of the PVData C++ package [14]. In PvaPy, the base class for all PV data types is *PvObject*, which represents a generic PV Structure. PvObject is initialized with a Python dictionary of PV introspection data, a set of key/value pairs describing the underlying PV structure in terms of field names and their types. The dictionary key is a string (the PV field name), and the value can be one of:

- PVTYPE: a scalar type, any of BOOLEAN, BYTE, UBYTE, SHORT, USHORT, INT, UINT, LONG, ULONG, FLOAT, DOUBLE, or STRING
- [PVTYPE]: a single element list, representing a scalar array

* Argonne National Laboratory's work was supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under contract DE-AC02-06CH11357.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

HIGH LEVEL APPLICATIONS FOR HLS-II

K. Xuan#, G. Liu, C. Li, J. Wang, L. Wang, W. Li, J. Y. Li
National Synchrotron Radiation Laboratory, University of Science and Technology of China, Hefei, Anhui 230029, China

Abstract

The Hefei light source was overhauled beginning from 2010 and completed in the end of 2013. The new light source is renamed as HLS-II. A set of high level application tools, including physical quantity based control IOC, lattice calibration tools, orbit feedback, etc., were developed for the light source commissioning and operation. These tools have been playing important roles in the commissioning and operation of the light source. This paper reports some critical applications.

INTRODUCTION

The Hefei light source was overhauled between 2010 and 2013 using a completely new structure. The new light source is renamed as HLS-II. The HLS-II is comprised of an 800 MeV Linac and an 800 MeV storage ring. The installation of the light source components were finished in the end of 2013. The machine commissioning was started in the beginning of 2014. In January 2015, the HLS-II was brought into operation. In order to make the commissioning more effective, and achieve high performance operation, a set of high level application tools were developed based upon the Experimental Physics and Industrial Control System (EPICS) [1]. These tools have been playing important roles in both the commissioning and the operation of the HLS-II. This paper reports some of these tools and their applications.

PHYSICAL QUANTITY BASED CONTROL SYSTEM

The physical quantity based control system was first developed at the Duke FEL laboratory [2]. With necessary modifications, the HLS-II control system is developed using familiar software structure as the Duke FEL control system. The functional diagram of the control system is illustrated in Fig. 1. The HLS-II control system is comprised of three parts, IOCs for hardware controls, soft IOCs for conversions between physical and engineering quantities, and OPIs for displaying essential parameters, and running high-level tools and a number of simulation programs.

The conversion between physical and engineering quantities is realized by a soft IOC. In many cases in the light source commissioning and operation, a set of related quantities (records) need to be simultaneously processed to keep the electron beam undisturbed. This requirement is fulfilled using the EVENT mechanism of EPICS. For example, when the energy (ENG) of the storage ring is changed, a set of event with designated values are posted, and all the records related to the focusing strength of quadrupole magnets are processed as they receiving these events. This scheme has been well serving for the energy ramping and lattice adjustment of the storage ring.

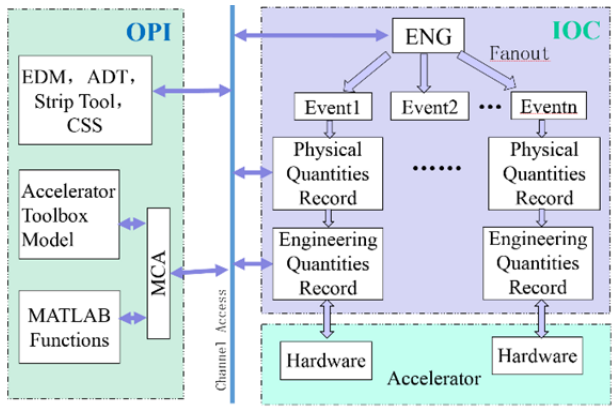


Figure 1: Functional diagram of the HLS-II control system.

The physical quantity of each component of the light source is usually divided into several parts for different purposes. For example, the total strength of a quadrupole magnet is given by

$$K = K_{design} + K_{adj} + K_{tune} + K_{comp}, \quad (1)$$

where K_{design} is the design value of that magnet, K_{adj} is used for adjustment, K_{tune} is used to change the transverse tunes, K_{comp} is used for lattice compensations. Another example is the kicking angle of an orbit corrector, its total kicking angle is composed of

$$\theta = \theta_{adj} + \theta_{feedback} + \theta_{comp}, \quad (2)$$

where θ_{adj} is used for adjustment, $\theta_{feedback}$ is used by the orbit feedback correction, and θ_{comp} is used to compensate residual field of some magnetic elements, such as undulators.

The physical quantity based control system of the HLS-II directly controls the physical quantities of the accelerator elements, including the magnetic field of various magnets, beam energy, and so on. This feature of the control system enables it to directly control the parameters of the accelerators and the electron beam. The physical quantities and related engineer quantities are automatically converted inside of the control system, and can be shared by different high level applications. This significantly improves the feasibility and effectiveness of the system. This system have been well serving the

xuanke@ustc.edu.cn

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

USE OF TORNADO IN KAT-7 AND MeeRKAT FRAMEWORK

C. De Villiers#, B. Xaia‡, SKA South Africa, National Research Foundation of South Africa, South Africa

Abstract

The KAT-7 and MeerKAT radio telescope control systems (www.ska.ac.za) are built on a rich Python architecture. At its core, we use KATCP (Karoo Array Telescope Communications Protocol), a text-based protocol that has served the projects very well. KATCP is supported by every device and connected software component in the system. However, its original implementation relied on threads to support asynchronous operations, and this has sometimes complicated the evolution of the software. Since MeerKAT (with 64 dishes) will be much larger and more complex than KAT-7, the Control and Monitoring (CAM) team investigated some alternatives to classical threading. We have adopted Tornado (www.tornadoweb.org) as the asynchronous engine for KATCP. Tornado, popular for Web applications, is built on a robust and very efficient coroutine paradigm that in turn is based on Python's generators. Co-routines avoid the complexity of thread re-entrancy and lifetime management, resulting in cleaner and more maintainable user code.

This poster will describe our migration to a Tornado co-routine architecture, highlighting the benefits and some of the pitfalls and implementation challenges we have met.

KATCP IN THE KAT-7 AND MEERKAT SYSTEMS

KATCP [1,2] is a simple ASCII communication protocol layered on top of TCP/IP.

It has been developed as a part of the Karoo Array Telescope (KAT) and MeerKAT projects and used at SKA South Africa for the monitoring and control of hardware devices. In this role it has been very successful and the specification is currently at Revision 5.

The original KATCP implementation provided a blocking client and a non-blocking CallbackClient.

Base message types are Request, Reply and Inform - the latter are sent asynchronously by a server to provide out-of-band data or (in some cases) to provide a way of segmenting the results of a previous request.

KATCP additionally defines a software Sensor type. Sensors are created with names and data types such as float, string etc. Dynamically a sensor may acquire a value and a status which it communicates to its registered listeners via callbacks. A listening client may set a strategy on the sensor which causes it to push its value and status to the listener periodically or on certain events, such as a value change. An ad-hoc query mechanism is also available.

KATCP messages, sensors, servers and clients are the building-blocks of the KAT-7 and MeerKAT Control and Monitoring systems. These robust abstractions support the next layer of the architecture, which comprises software proxies to abstract access to real hardware, and components that partition the work of system startup and shutdown, scheduling, observation control and monitoring.

LIMITATIONS OF THREADS

In earlier implementations of KATCP, the inherent concurrency of real-time processes was modelled using software thread. Threads provide the illusion of parallel execution within a single processor core by performing pre-emptive task switches at the system level. This has benefits for processing efficiency since the system need never be idle - while some thread of control is awaiting an event, another thread can be executing. This can be especially useful in networked systems where many cycles would otherwise be wasted waiting for I/O events.

However threads also have some well-known drawbacks. Each thread has its own execution context and this means that threading is resource-intensive, so that the number of active threads must be limited. Perhaps even more important is the complexity they introduce into software design. Since any thread of control may be interrupted or resumed at essentially arbitrary moments, software becomes ‘non-linear’ and the designer must carefully guard against inadvertent corruption of shared resources. Numerous best-practices and software constructs exist to alleviate these problems, but all contribute to the complexity and cost of software development and maintenance.

Finally the Python language, which has proved immensely valuable in the development of our systems, implements a Global Interpreter Lock (GIL) which essentially disables threading for compute-bound tasks on a single processor.

MIGRATION TO TORNADO

Like many teams facing these challenges, we have been interested in the developments in coroutine-based concurrency frameworks. Coroutines are a generalisation of subroutines based on co-operative multitasking, in contrast to the pre-emptive model used by threads Coroutines differ from subroutines in allowing multiple entry-points (and multiple entries per entry-point) within the body of a routine, with the preservation of the full execution context at that point. Because task-switching is

#charles@ska.ac.za
‡ bxaia@ska.ac.za

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

BUNCH TO BUCKET TRANSFER SYSTEM FOR FAIR

J.Bai[1,2], T.Ferrand[1,3], D.Beck[1], R.Bär[1], O.Kester[1,2], D.Ondreka[1], C.Prados[1],W.Terpstra[1]
1. GSI Helmholtzzentrum für Schwerionenforschung, Darmstadt, Germany
2. IAP, Goethe Universität Frankfurt, Frankfurt am Main, Germany
3. Technische Universität Darmstadt, Darmstadt, Germany

Abstract

For the FAIR accelerator complex, synchronization of the bunch to bucket (B2B) transfer will be realized by the General Machine Timing system and the Low-Level RF system. Based on these two systems, both synchronization methods, the phase shift and the frequency beating method, are available for the B2B transfer system for FAIR. This system is capable to realize the B2B transfer within 10ms and the precision better than 1° for ions over the whole range of stable isotopes. At first, this system will be used for the transfer from the SIS18 to the SIS100. It will then be extended to all transfers at the FAIR accelerator facility. This paper introduces the synchronization methods and concentrates on the standard procedures and the functional blocks of the B2B transfer system.

INTRODUCTION

FAIR is aiming at providing high-energy beams with high intensities. Based on the existing GSI accelerators UNILAC and the SIS18 serving as an injector, high intensity ion beams over the whole range of stable isotopes will be accelerated in the new heavy ion machine SIS100 to higher energies. The FAIR new accelerator complex with storage rings will in the full version consist of the SIS100, the Collector Ring CR, the accumulator/decelerator ring RESR and the New Experimental Storage Ring NESR. An additional High Energy Storage Ring HESR serves experiments with high energy antiprotons. The B2B transfer means that one bunch of particles, circulating inside the source synchrotron, is transferred into the center of a bucket of the target synchrotron. For the FAIR project, there are many transfers involving the bunch to bucket (B2B) transfer. E.g. the B2B transfer from the SIS18 to the SIS100, from the SIS18 to the ESR, from the SIS100 to the CR, from the CR to the HESR and later to RESR.

For the FAIR accelerator complex, synchronization of the B2B transfer system will be realized by General Machine Timing (GMT) system [1] and Low-Level RF (LLRF) system) [2].

The main tasks of the GMT system are time synchronization of more than 2000 nodes with nanosecond accuracy, distribution of timing messages and subsequent generation of real-time actions by the nodes of the timing system. The GMT consists of a Data Master (DM), a Clock Master (CM), the White Rabbit (WR) timing network and integrates nodes [3]. The DM schedules actions by broadcasting messages, which will be received and executed by the corresponding node at the designated time.

For the synchronization of LLRF system, the GMT system is complemented and linked to Bunchphase Timing System (BuTiS) [4]. BuTiS is a campus wide clock synchronization and distribution system, locally generating two delay compensated high precision clock signals with a jitter of 10 femtosecond.

Dependent on the BuTiS, the B2B transfer system at FAIR is distinguished from others. For CERN and other accelerator facilities, the B2B transfer system is based on the cavity measured signals [5]. Instead for FAIR, it is based on the driven signals directly derived from the BuTiS.

METHODS OF SYNCHRONIZATION TWO SYNCHROTRONS

For the proper transfer, the phase advance between the bunch of the source synchrotron and the bucket of the target synchrotron must be precisely controlled before the bunch is ejected. The process of achieving the detailed phase adjustment is usually termed "synchronization". There are usually two methods available for the synchronization process, the phase shift method and the frequency beating method.

A short description of the phase shift method is as follows. At a scheduled time well before ejection, the phase of the beam in the source synchrotron and the phase of the bucket in the target synchrotron are measured with respect to the phase of a common central reference signal, which is synchronously distributed to the source and target synchrotrons. Based on the measured phase, the radio frequency (RF) frequency of the source or target or both synchrotrons is modulated away from the nominal value for a period of time and then modulated back so that the phase shift created by the frequency modulation could compensate for the expected phase difference. After the phase shift, the bunches can be injected into buckets. If the RF frequency of the source and target synchrotrons are same, it brings an infinite time frame in ideal situation beginning from the end of the phase shift process, within which the bunches can be injected into buckets at any time. This is the so called synchronization window for the phase shift method (See Figure 1). The phase shift process must be performed adiabatically for the longitudinal emittance to be preserved [6]. But when the target synchrotron has no particles, the phase shift can be done for the target synchrotron without adiabatical consideration (e.g. Phase jump).

The frequency beating method uses the effect of two RF signals of slightly different frequencies, perceived as periodic variations in phase difference whose rate is the difference between the two frequencies. The RF frequency of the

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

TIMING SYSTEM AT MAX IV - STATUS AND DEVELOPMENT

J. Jamroz, V. Hardion, V. Martos, J. Forsberg, D. Spruce, MAX-IV, Lund, Sweden

Abstract

A MAX IV construction of two storage rings (SR1.5GeV and SR3GeV) and a short pulse facility (SPF) has been proceeding over last years and will be finished in the middle of 2016. In 2014, few timing procurements were successfully finalized according to the MAX IV requirements (see [1] for details) and the installation works are ongoing along with the TANGO control system integration. The design covers the timing synchronization (TIM) and acquisition, fast orbit feedback (FOFB), fast machine protection system (FMPS) and integration of MAX IV operation modes. The LINAC commissioning started in 2014 and was successful together with the first beam line (FEMTOMAX). The SR3GeV commissioning started in August 2015 and is ongoing. The SR1.5GeV is being installed and will be commissioned in 2016.

MAX IV OPERATION MODES

MAX IV will work in 4 different modes:

- LINAC 3GeV (LIN)
- LINAC 3GeV + Short Pulse Facility (SPF)
- LINAC 1.5GeV + Storage Ring 1.5GeV (SR1)
- LINAC 3GeV + Storage Ring 3GeV (SR3)

LIN and SPF are rated for 100Hz injection whereas SR1 and SR3 for 10Hz.

Each mode is synchronized to its own radio frequency (RF). LIN and SPF are driven by 39 harmonics of the LINAC master oscillator (MO) 3GHz (2998500000/39=76884615.38Hz in reality, around 77MHz). The 39 harmonic is required for a synchronization of two laser systems, one at the beginning of the LINAC, other 400m away in the short pulse facility. This implementation also keeps the same constant phase for other frequencies: 100Hz LINAC injection trigger, 1kHz laser pumping frequency etc.

SR1 is locked to 100MHz (RF1) and SR3 is locked to 100MHz (RF3) but those frequencies are a bit different as the real value comes from the construction and temperature parameters of storage rings. Both rings work in the same way, each SR TIM generates a LINAC injection trigger which is locked to its RF and in constant phase to its machine clock (MC). The MC phase shift (10ns step) defines the SR bucket number. SR3 stores 176 electron buckets and its MC is equal to 568kHz, whereas SR1 stores 32 electron buckets and its MC is equal to 3.12MHz.

On top of that, the synchronization to mains electricity (ME) 50Hz will be implemented, so the injection will always be applied under constant parameters of high power devices.

In summary, 3 different LINAC triggers are delivered:

- LIN and SPF – TRIG0 (3GHz, 77MHz, 50Hz),
- SR1 – TRIG1 (100MHz, 3.12MHz, 50Hz),
- SR3 – TRIG3 (100MHz, 568kHz, 50Hz).

The brackets include a set of frequencies which are synchronized one other. The future timing improvement considers adding 3GHz master oscillator (MO) to TRIG1 and TRIG3 (see “MO 3GHz SYNCHRONIZATION” section for more information).

SWITCHING MODES

Switching over each operation mode has not been defined yet.

Main assumptions:

- The LINAC trigger runs 24/7 regardless other modes, so LINAC devices keep constant parameters (temperature etc.).
- Max LINAC death time (no trigger) around 5 second.
- To avoid an overcurrent injection to the nominal stored beam (500mA) in SR1 and SR3 without dumping the beam (BEAMDUMP).
- To avoid an overcurrent injection to a single storage ring bucket, so not to excess the nominal 5nC/bucket.
- Full automatic operation, one instruction/button of the operator should be enough to run the machine.
- Once SR1 and SR3 are “full”, SPF/LIN mode should be activated (nominal conditions).
- Except TIM, other devices (which are much slower) will be a part of the switching mechanism (e.g. SR1 and SR3 LINAC bending magnets etc.).
- Each timing system (SPF, SR1, SR3) has to operate autonomously 24/7 regardless the actually selected mode, delivering acquisition triggers, beam pinging, MC, FMPS functionality etc.

CORE HARDWARE

Beam Position Monitors (BPM-s)

The BPM system is based on Libera Brilliance+ which is a uTCA 2U box containing:

- 1x timing module (EVRX) - connected to the timing system via an optical link,
- 3x or 4x BPM processor unit - connected to RF pick ups in the storage rings,
- 1x gigabit data exchange module (GDX) which exchanges the beam orbit data/calculations,
- 1x RS485 card controlling power supplies which are connected to ring corrector magnets,
- 1x CPU module running Linux and connected to Ethernet.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

OPERATION STATUS OF J-PARC TIMING SYSTEM AND FUTURE PLAN

N. Kamikubota#, N. Kikuzawa, F.Tamura, and N. Yamamoto
J-PARC-Center, KEK & JAEA, Tokami-mura, Ibaraki, Japan

Abstract

The beam commissioning of J-PARC started in 2006. Since then, the timing system of J-PARC has contributed stable beam operation of accelerators. The present timing system is reviewed from the aspects of history, fiber-optic cable network, VME modules and their configurations, followed by upgrade studies for the future.

INTRODUCTION

J-PARC (Japan Proton Accelerator Research Complex) is a high-intensity proton accelerator complex, located in Ibaraki, Japan. It consists of three accelerators: a) 400-MeV linac (LI), b) 3-GeV Rapid Cycling Synchrotron (RCS), and 30-GeV Main Ring (MR), and three experimental facilities: d) Material and Life Science Experimental Facility (MLF), e) Neutrino Experimental Facility (NU), and f) Hadron Experimental Facility (HD) [1-3].

In J-PARC, there are two time cycles. The rapid cycle, 25 Hz, is used at LI, RCS and MLF. Through LI was designed to be operated at the 50-Hz repetition rate, current operation is 25 Hz. The slow cycle is used at MR, NU and HD. In 2015, when MR delivers proton beams to NU (HD), 2.48s (6.00s†) is used, respectively. Since the slow cycle determines the overall time behaviour of accelerators, it is often called "machine cycle". Two different cycles are co-exist in J-PARC.

The J-PARC Timing System started operation since 2006. The purpose is to provide a trigger to each of the accelerator components with a specified delay. In general, only one trigger is necessary within the 40ms (2.48s/6.00s) time slot for the rapid-cycle (slow-cycle) machines, respectively.

PRESENT TIMING SYSTEM

History and Overview

The early studies for the J-PARC controls were carried out in 2003 [4], in which the design for the timing system was one of the issues. In 2003, the VME-bus modules for timing control, both the send module and the receiver module were developed in collaboration with domestic companies [5]. Mass productions of VME modules and related NIM modules were in 2005-2007. The first beam to the LI (RCS, MR) was in 2006 (2007, 2008), respectively. Summarized history is shown in Figure 1.

#norihiko.kamikubota@kek.jp
† 5.52s after RUN64, October 2015

The control system for the J-PARC, including the timing system, is based on EPICS toolkit [6]. Both low-level support (device driver and EPICS databases) and high-level applications for the timing system were developed by ourselves. In order to manage table-formatted registers of VME modules, the waveform record-type of EPICS is used. Java and Python have been preferred for high-level GUI applications. .

In 2015, we have 118 (43, 45) receiver modules, which correspond to ~540 (~220, ~300) end-point signals for LI (RCS, MR), respectively. In addition, a few receiver modules exist for each of three experimental facilities. Only one send module exists for the whole accelerator complex.

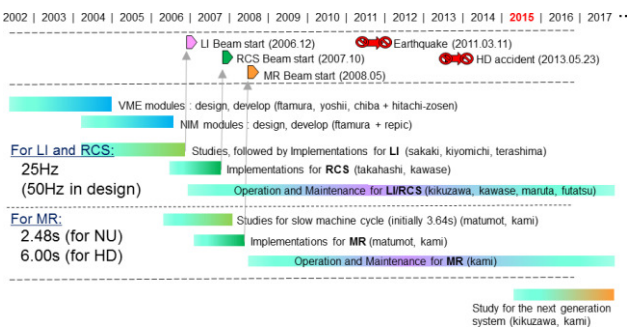


Figure 1: History of the J-PARC timing system.

Base-signals and Their Distribution

We have three "base-signals": a) 12MHz master clock (CLK), b) 25Hz trigger (Trig), and c) type-code (Type). The CLK is provided by a commercial high-stability function generator. The Trig is generated from the master clock, and used as the start signal of each rapid-cycle. The Type, a 32-bit number, is sent at the Trig rate after serialized. A receiver module uses it to generate delayed trigger signals during the next rapid-cycle.

We have fiber-optic cable networks throughout our facility buildings. Three base-signals are generated in the Central Control Building (CCB), then distributed to all the facility buildings using the fiber network (Figure 2).

A schematic view of base-signal distribution is shown in Figure 3. Several E/O, O/E and fan-out modules are used. All of the receiver modules in facility buildings receive the same base-signals.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

REAL-TIME PERFORMANCE IMPROVEMENTS AND CONSIDERATION OF PARALLEL PROCESSING FOR BEAM SYNCHRONOUS ACQUISITION (BSA)*

K. H. Kim#, S. Allison, T. Straumann, E. Williams
SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

Abstract

Beam Synchronous Acquisition (BSA) provides a common infrastructure for aligning data to each individual beam pulse, as required by the Linac Coherent Light Source (LCLS)[1]. BSA allows 20 independent acquisitions simultaneously for the entire LCLS facility and is used extensively for beam physics, machine diagnostics and operation. BSA is designed as a part of LCLS timing system [2,3] and is currently an EPICS record based implementation, allowing timing receiver EPICS applications to easily add BSA functionality to their own record processing. However the lack of real-time performance of EPICS [4] record processing and the increasing number of BSA devices has brought real-time performance issues. The major reason for the performance problem is due to the lack of separation between time-critical BSA upstream processing and non-critical downstream processing. We are improving BSA with thread level programming, breaking the global lock in each BSA device, adding a queue between upstream and downstream processing, and moving out the non-critical downstream to a lower priority worker thread. We are also investigating the use of multiple worker threads for parallel processing in Symmetric Multi-Processor (SMP) system.

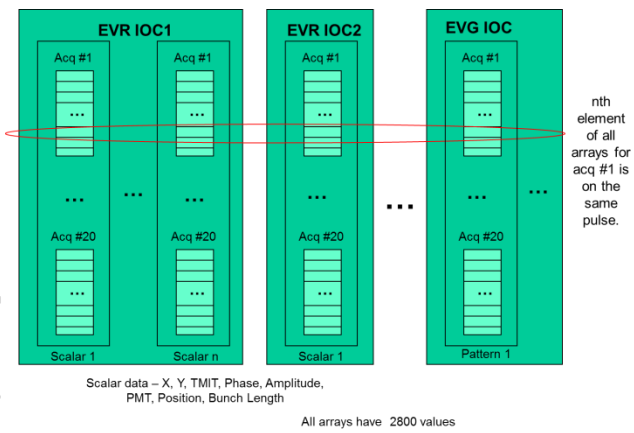


Figure 1: Data Acquisition across IOCs

BEAM SYNCHRONOUS ACQUISITION

BSA has been designed as a part of event system in LCLS-I to acquire all of beam dependent scalars across multiple IOCs in the entire accelerator facility on the same beam pulse, allowing correlation analysis using the pulse by pulsed aligned acquisition data (Figure 1). BSA acquires up to 2,800 values per scalar in one acquisition request; each value of the 2,800 can be an average of up to 1,000 values. It also provides RMS values and other statistics. The BSA can process 20 different acquisitions simultaneously [5].

BSA is implemented in three parts. A user request for an acquisition is done by EPICS CA client. Data gathering is processed on the Event Generator (EVG) and Event Receiver (EVR) IOCs. When gathering is finished, access of prepared data waiting on IOCs is done by CA clients, with checks for a good acquisition.

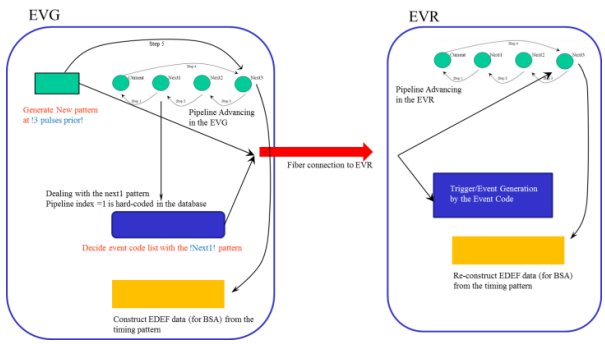


Figure 2: 360Hz tasks in EVG and EVR

360Hz Task in EVG

Acquisition setup and start requests done on the EVG IOC. The EVG IOC performs 360Hz checking and user notification when a BSA is finished. A 360Hz event task wakes up on an interrupt from a clock synched AC powerline zero-crossing which also provides timeslot for the event system. The event task is the heart of EVG IOC, generating a timing pattern for 3 pulses ahead and schedules timing events for the next pulse. The event task also checks for a match between the new pulse's timing pattern, beam code, and each active acquisition for BSA. It keeps a count of the number of measurements and the number of values in the current average per acquisition. One part of the timing pattern represents which acquisitions are matched. The timing pattern is broadcasted to EVRs and contains pulse id, timestamp, and additional BSA information. The pulse information is pipelined, thus the timing pattern is for 3 pulses ahead (Figure 2).

*Work supported by the the U.S. Department of Energy, Office of Science under Contract DE-AC02-76SF00515 for LCLS I and LCLS II.
#khkim@slac.stanford.edu

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

APPLICATION USING TIMING SYSTEM OF RAON ACCELERATOR

S. Lee*, CW. Son†, HJ. Jang‡, IBS, Daejeon, South Korea

Abstract

RAON is a particle accelerator to research the interaction between the nucleus forming a rare isotope as Korean heavy-ion accelerator. RAON accelerator consists of a number of facilities and equipments as a large-scaled experimental device operating under the distributed environment. For synchronization control between these experimental devices, timing system of the RAON uses the VME-based EVG/EVR system. This paper is intended to test high-speed device control with timing event signal. To test the high-speed performance of the control logic with the minimized event signal delay, we are planing to establish the step motor controller testbed applying the FPGA chip. The testbed controller will be configured with Zynq 7000 series of Xilinx FPGA chip. Zynq as SoC (System on Chip) is divided into PS (Processing System) and PL (Programmable Logic). PS with the dual-core ARM cpu is performing the high-level control logic at run-time on linux operating system. PL with the low-level FPGA I/O signal interfaces with the step motor controller directly with the event signal received from timing system.

This paper describes the content and performance evaluation obtaining from the step motor control through the various synchronized event signal received from the timing system.

INTRODUCTION

The RAON [1] is a new heavy ion accelerator under construction in South Korea, which is to produce a variety of stable ion and rare isotope beams to support various researches for the basic science and applied research applications. To produce the isotopes to fulfill the requirements we have planed the several modes of operation scheme which require fine-tuned synchronous controls, asynchronous controls, or both among the accelerator complexes. RAON, which is a large experimental machine, consists of many experimental devices and additional facilities. For synchronized control under the distributed environment, timing system of the RAON uses the VME-based Event Generator (EVG)/Event Receiver (EVR) [2] system.

TIMING SYSTEM

The timing system uses the EVG/EVR system of the Micro-Research Finland Oy company in VME form factor. The timing system consists of an EVG which converts timing events to optical signals, Fan-Out unit that repeats the signals, an array of EVRs and VME-controller. The EVRs decode the optical signal and produce hardware signal or software interrupt for the timing event. It uses the MVME6100 or 3100 of Emerson company as VME-controller. The controller also uses real-time operating system, VxWorks 6.9 for a fast response. Finally, EPICS IOC (mrfioc2) on VxWorks controls the FPGA-based EVG/EVR boards.

The characteristics of MRF timing system are :

- Event driven system, 256 event codes
- Event generation using external RF reference clock
- 50 ~ 125MHz event clock rate
- Events generated
  - From external HW inputs
  - Two sequencers (up to 2048 events/sequencer)
  - Multi counters
- Cascaded Event Generators
- Different Clock Synchronization

Hardware Configuration

Hardware configuration of the timing system consist of:

- XLI GPS Time System
- Rubidium Frequency Standard Clock Source (FS725)
- Event Trigger System (EVG/EVR,Fanout Repeater/Concentrator)
- MVME 6100/MVME 3100 controller board
- VME Crate (Wiener)
- SMA 100a RF Signal Generator

Figure 1 shows hardware configuration to test timing system.

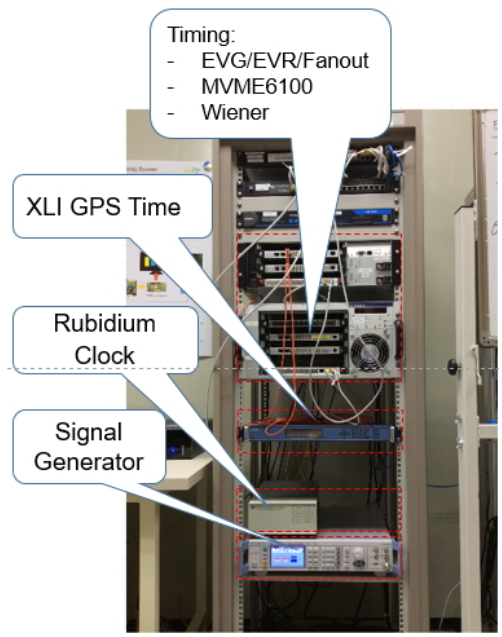


Figure 1: Timing hardware configuration.

* silee7103@ibs.re.kr
† scwook@ibs.re.kr
‡ lkcom@ibs.re.kr

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

PROTOTYPE OF WHITE RABBIT NETWORK IN LHAASO

HongMing Li*, GuangHua Gong*, Tsinghua University, Beijing, China
Qiand Du, LBNL, USA

Abstract

Synchronization is a crucial concern in distributed measurement and control systems. White Rabbit provides sub-nanosecond accuracy and picoseconds precision for large distributed systems. In the Large High Altitude Air Shower Observatory project, to guarantee the angular resolution of reconstructed air shower event, a 500 ps overall synchronization precision must be achieved among thousands of detectors. A small prototype built at Yangbajin, Tibet, China has been working well for a whole year. A portable calibration node directly synced with the grandmaster switch and a simple detectors stack named Telescope are used to verify the overall synchronization precision of the whole prototype. The preliminary experiment results show that the long term synchronization of the White-Rabbit network is promising and 500 ps overall synchronization precision is achievable with node by node calibration and temperature correction.

INTRODUCTION

The Large High Altitude Air Shower Observatory (LHAASO) project is a dedicated instrument for searching the origin of galactic cosmic rays, which consists of 4 sub-detector arrays. The 1km^2 complex array (LHAASO-KM2A) includes 5635 scintillation electron detectors and 1221 muon detectors. 500 ps (rms) overall synchronization precision must be achieved among the spread thousands of detectors to guarantee the angular resolution of reconstructed air shower event. [1][2] A small prototype based on the White Rabbit (WR) network [3][4] is built at Yangbajin, Tibet, China on August, 2014. This paper talks about the deployment and synchronization monitor of the prototype.

DEPLOYMENT

Components and Topology

This prototype contains 4 WR switches and 50 WR customized nodes (48 electron detectors and 2 muon detectors). Rubidium clock constrained GPS receiver takes the clock and frequency source of the whole network. The four WR switches is used to build a four layer hierarchy WR network. The grandmaster switch (GMS) uses the frequency from GPS receiver and synchronizes its time (seconds counter) through NTP protocol from GPS receiver. Each WR node synchronizes its local clock against the GMS through fibres of one hundred meters and uses the synchronized clock to timestamp the events. Detail information of each component is listed below:

- 8040C, Rubidium Frequency Standard
- XL-GPS, produced by symmetricom
- WR switch, version V4.0.1
- Compact Universal Timing Endpoint (CUTE) [5], customized WR node, WRPC version V2.1. [6]

Locations

The prototype is deployed in the ARGO experiment hall whose altitude is around 4300m. The locations of WR switches and nodes are showed in Figure 1.

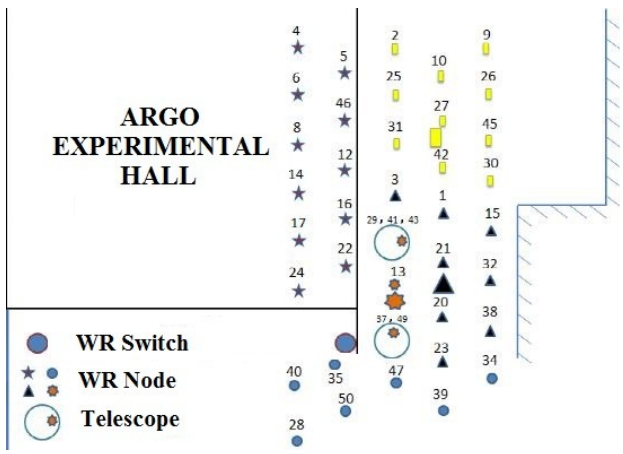


Figure 1: Locations of WR Switches and Nodes

Rubidium Frequency Standard, GPS, GMS and host computers are put in the ARGO Experiment Hall with small temperature variations, while the other three switches are put together in a waterproof steel box out of the Experiment Hall. A normal switch is used to connect the GPS network interface (acts as NTP server), host computer and all four WR switches for remote control.

The telescope is a pyramid of several WR nodes for monitoring the synchronization between the nodes. One node is on the top of another node to make sure they detect the same signal simultaneously.

Calibration Procedure

By following the calibration procedure given in the document [7], a node by node calibration is applied on all WR devices to conquer the differences caused by the components diversity.

As the bad control of the power noise in these boards, the synchronization precision (standard deviation value of the Pulse Per Second (PPS) skew) of each link is decreased to 100ps. Then the synchronization accuracy (absolute mean value of the PPS skew) less than 200 ps after calibration is accepted.

* The authors are with Key Laboratory of Particle and Radiation Imaging, Tsinghua unversity, Beijing, China.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

A GENERIC TIMING SOFTWARE FOR FAST PULSED MAGNET SYSTEMS AT CERN

C. Chanavat, M. Arruat, E. Carlier, N. Magnin, CERN, Geneva, Switzerland

Abstract

At CERN, fast pulsed magnet (kicker) systems are used to inject, extract, dump and excite beams. Depending on their operational functionalities and as a result of the evolution of controls solutions over time, the timing controls of these systems are based on different hardware architectures that result in a large disparity of software solutions. A Kicker Timing Software (KiTS), based on a modular hardware and software architecture, has been developed with the objective to increase the homogeneity of fast and slow timing control for fast pulsed magnet systems. The KiTS uses a hardware abstraction layer and a configurable software model implemented within the Front-End Software Architecture (FESA) framework. It has been successfully deployed in the control systems of the LHC and SPS injection kickers, the SPS extraction kickers and the SPS tune measurement kickers.

INTRODUCTION

A kicker system must meet two requirements that are contradictory i.e. a high deflection strength and a short rise time of its magnetic field. Both properties are, for a given operational voltage, proportional to the product of the electrical current passing through the magnet and its length.

In order to meet these contradictory requirements, the length of the magnet is reduced by splitting the system into several independently powered magnet modules. Generally, the powering circuit of a magnet module consists of a set of well identified hardware components:

- A resonant charging power supply (RCPS) used to charge a pulse generator in order to decrease the time interval between two successive pulses and reduce the number of faulty shots;
- A line type pulse generator based on a pulse forming network (PFN) supplying quasi rectangular current pulses with variable length and amplitude;
- Up to three high voltage fast switches used to transfer in a controlled way the energy stored in the PFN to the magnet and to adjust the pulse length;
- A coaxial transmission line (TL) connecting the generator to ferrite type magnets;
- A ferrite magnet, built as lumped parameter delay lines and working in ultra-high vacuum;
- A termination resistor (TR) matched to the characteristics impedance of the generator, transmission line and magnet and absorbing the pulse energy supplied by the generator.

On this basis, the simplest kicker system comprises at least one PFN charged by one RCPS and discharged through the TL within one terminated magnet by one high voltage switch as shown in Figure 1.



Figure 1: Simplest kicker system architecture

Although peak current, system impedance, pulse shape, pulse duration and repetition rate are different for every system, the different hardware components have been standardised as far as possible, and, depending of the kicker functions (beam injection, extraction or excitation), the number of RCPS, PFN, high voltage switches and magnet are combined in more or less complex architecture in order to provide the required functionalities. Different types of combination used in the SPS and the LHC are summarised in Table 1.

Table 1: Example of Kicker Architecture

System	RCPS per System	PFN per RCPS	Switch per PFN	Magnet per PFN
SPS Injection	4	2	3	2
SPS East Extraction	1	5	2	1
SPS West Extraction	1	1	1	4
SPS Tune	4	1	3	1
LHC Injection	2	2	2	1

Up to now the timing control of these different hardware combinations was based on dedicated real-time software strongly linked not only to the kicker hardware configurations themselves but also to the beam process where the system was used. With time, this approach has resulted in a high number of software packages to be maintained and to a high dependency of each software package with machine operation conditions.

Additionally, as this approach has been used for more than 40 years, a high diversity of electronic modules is used to generate the delays needed to trigger the different high voltage switches. As the access to these different timing delays is strongly embedded inside the real-time software, maintenance is now becoming more and more difficult due to the obsolescence of the delay modules and the difficulties to replace them without having to do a full re-engineering of the actual software.

In order to solve these two problems, a generic kicker timing software constructed on the basis of the existing set of standardised hardware components has been

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

STATUS DEVELOPMENT OF SIRIUS TIMING SYSTEM

J.L.N. Brito*, S.R. Marques, D.O. Tavares, L.A. Martins, LNLS, Campinas, Brazil

Abstract

Sirius is a new low-emittance 3 GeV synchrotron light source under construction in Brazil by LNLS, scheduled for commissioning in 2018. Its timing system will be responsible for providing low jitter synchronized signals for the beam injection process as well as reference clocks and triggers for diverse subsystems such as electron BPMs, fast orbit feedback and beamlines distributed around the 518 meters circumference of the storage ring, Booster and Linac. It will be composed of Ethernet-configured standalone event generators and event receivers modules developed by SINAP through a collaboration with LNLS. The modules will be controlled by remote EPICS soft IOCs. This paper presents the system structure and the status of the development, some options for integrating it to the Sirius BPM MicroTCA platform are also discussed.

INTRODUCTION

Sirius light source consists of a 150 MeV Linac, a 150 MeV to 3 GeV booster synchrotron and a 3 GeV, 518 meters circumference, storage ring with 20 straight sections, achieving a very low beam emittance of 0.24 nm·rad [1]. Many devices must be triggered synchronously by a low jitter timing system to guide the electrons from the Linac to the storage ring.

The main purpose of a timing system is to generate and distribute deterministic signals to control the beam injection process. Secondary purposes are to provide timestamps for diverse subsystems to ensure a time-coherent behaviour of data acquisition and to deliver reference clock for the electron BPM and fast orbit feedback systems.

The range of possibilities considered to perform the Sirius timing system can be summarized in four major approaches: (i) a completely commercial timing system [2], (ii) a solution entirely based on scientific test and measurement instruments [3], (iii) an in-house development and (iv) collaborative options, such as White Rabbit Project [4] and SINAP timing system [5]. After surveying these possibilities, the SINAP timing system was chosen based on the available functionality it could offer, proven reliability [6] and collaboration opportunity.

REQUIREMENTS

The Sirius injection system operates with a 2 Hz repetition rate, which must be synchronized to the mains at 60 Hz to turn repetitive the ripple effect of the power supplies on the injected beam in the accelerators. Trigger signals should be derived from the RF frequency to synchronize pulsed elements on Linac, Booster and storage ring to electron bunches. The 90 keV e-gun, that will operate in single-bunch and multi-bunch modes, is the pulsed element with the most demanding trigger signal, requiring jitter less than 50 ps rms, coarse delay resolution of 2 ns and fine delay resolution of 20 ps.

After the commissioning phase, top-up is foreseen to be the Sirius normal operation mode, in which a bunch current measurement system to take the bunch current profile may be necessary to perform the bunch selection. Hence, Sirius timing system should support injection to any bunch and software integration to external systems. Furthermore, other requirements include that critical signals needs to be transmitted through optical fibres to minimize EMI, electrical outputs requires TTL level compatibility and inputs for external triggers are essential.

Table 1 presents the main parameters about Sirius and timing system specifications.

Table 1: Main Sirius Parameters and Timing Specifications

Parameter	**Value**
RF freq. (Storage Ring & Booster)	499.658 MHz
Storage ring circumference	518.4 m
Storage ring harmonic number	864
Storage ring revolution freq.	0.578 MHz
Booster circumference	496.8 m
Booster harmonic number	828
Booster revolution freq.	0.603 MHz
Coincidence number	19872
Coincidence frequency	25.144 kHz
Linac RF frequency	2.997 GHz
Repetition rate	2 Hz
Specification	**Value**
E-gun rms jitter	< 50 ps
E-gun trigger coarse delay step	2 ns
E-gun trigger fine delay step	20 ps
General trigger delay step	8 ns
Electrical outputs	TTL level
Event clock freq. ($\frac{RF\ freq.}{4}$)	124.915 MHz

SYSTEM DESIGN

Sirius timing system is event-based [7], i.e. an event generator (EVG) broadcasts event frames to event receivers (EVR and EVE) through a star topology optical fibre network. An event frame consists of an 8-bit event code and an 8-bit distributed data bus. Each decoded event can generate configurable output triggers and each bit of the distributed data bus maps a channel clock, which can be configured to any integer sub multiple of the event clock. Since the event clock is derived from the RF frequency, this and the channel clocks are synchronized. Therefore, an event

* joao.brito@lnls.br

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

CERN TIMING ON PXI AND cRIO PLATFORMS

A. Rijllart, O. O. Andreassen, J. Blanco Alonso, CERN, Geneva, Switzerland

Abstract

Given the time critical applications, the use of PXI and cRIO platforms in the accelerator complex at CERN, require the integration into the CERN timing system. In this paper the present state of integration of both PXI and cRIO platforms in the present General Machine Timing system and the White Rabbit Timing system, which is its successor, is described. PXI is used for LHC collimator control and for the new generation of control systems for the kicker magnets on all CERN accelerators. The cRIO platform is being introduced for transient recording on the CERN electricity distribution system and has potential for applications in other domains, because of its real-time OS, FPGA backbone and hot swap modules. The further development intended and what types of applications are most suitable for each platform will be discussed.

INTRODUCTION

At CERN, the accelerator operation requires nanosecond precise timing for many systems that act on the particle beams, such as for injection, acceleration, focusing, measurement and extraction. These systems are synchronised and triggered using the General Machine Timing (GMT) [1] system, which has been custom designed at CERN. GMT electronic boards have been developed for a several bus systems, mainly VME and cPCI. At the time of choice in 2004, PXI and cRIO were not mature enough [2], however, during the last 5 years their success in a large variety of industrial areas [3,4], together with the steady improvement of the graphical system design tools [5] have made the PXI and cRIO platforms an interesting candidate for accelerator systems. But no GMT electronic boards fit in the PXI and cRIO bus systems and the C++ driver software for Linux was not compatible with the operating systems used on these platforms.

The successor of the GMT system [6], called White Rabbit, provides improved timing accuracy and solves most of the GMT's shortcomings, and is actively being developed at CERN. An IEEE standardisation committee has been set up to incorporate it into a new IEEE-1588 standard.

In this paper we describe the adaptation of the GMT to the PXI, the White Rabbit timing on the cRIO and how we envision the White Rabbit on the PXI, to enable CERN engineers and physicists to profit from advantages of industrial platform with CERN timing features.

GMT ON PXI

The CERN accelerator control standards for front-end systems are the VME and cPCI busses. Several equipment groups have chosen PXI, because of their instrumentation needs. However, to be able to use the CERN timing (GMT) a VME or cPCI had to be added to house the timing receiver and run the timing software. The generated triggers were then wired into the PXI either by copper or optical links. When one timing receiver could serve many PXI systems, such as for the collimator control [7], this small overhead is acceptable. For systems with a few PXI crates it would be an advantage if the timing card could be housed in the PXI crate.

For the hardware we use a GMT timing receiver card in PMC format [8] developed at CERN and a PMC carrier for PXI [9] from industry. This carrier is actually a cPCI carrier and therefore compatible with standard PXI crates and PXI Express crates with hybrid slots (PXI/PXIe) (see slot 5 in Fig. 1).

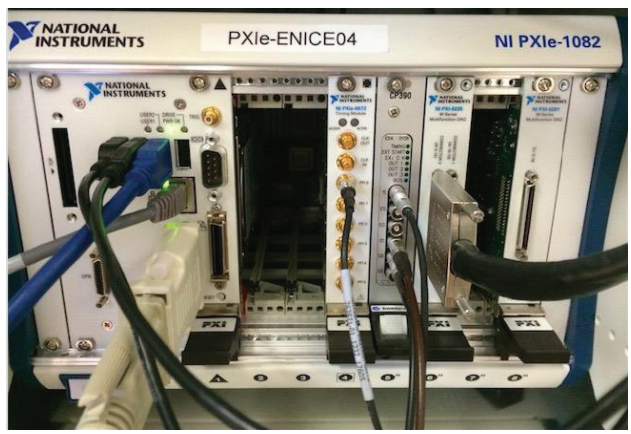


Figure 1: The GMT module installed in slot 5 of a PXI Express crate.

For the software we run a Hypervisor [10] on the PXI controller. This enabled us to run the existing timing library on Scientific Linux 6 on one core of the CPU and the LabVIEW RT system (Pharlap) on the other. In such a system the triggers generated by the GMT receiver were wired to a PXI trigger module that could provide bus triggers for commercial data acquisition cards in the rest of the PXI crate. The Hypervisor solution has also be used in other laboratories for the same purpose, such as in the Brazilian Synchrotron Light Laboratory (LNLS) [11] have used the Hypervisor to profit from the best features of two different operating systems.

WHITE RABBIT ON cRIO

The White Rabbit (WR) node is an open hardware design, accessible from CERN's open hardware repository [12]. Engineers at University of Zürich have used the WR design to build a C-series module for the cRIO platform [13], using the Module Development Kit of National Instruments. They adapted the module on FPGA code and wrote a driver in LabVIEW FPGA for the control and readout of the WR time stamp.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

AN UPDATE ON CAFE, A C++ CHANNEL ACCESS CLIENT LIBRARY, AND ITS SCRIPTING LANGUAGE EXTENSIONS

J. Chrin, Paul Scherrer Institut, 5232 Villigen PSI, Switzerland

Abstract

CAFE (Channel Access interFacE) is a C++ client library that offers a comprehensive and easy-to-use interface to EPICS (Experimental Physics and Industrial Control System). Functionality is provided for the synchronous and asynchronous interaction of individual and groups of low-level control data, coupled with an abstraction layer to facilitate development of high-level applications. The code base has undergone major refactoring to make the internal structure more comprehensible and easier to interpret, and further interfaces have been implemented to increase its flexibility, in readiness to serve as the CA host in fourth-generation and scripting languages for use at SwissFEL, Switzerland's X-ray Free-Electron Laser facility. An overview of the structure of the code is presented, together with an account of newly created bindings for the Cython programming language, which offers a major performance improvement to Python developers, and an update on the CAFE MATLAB Executable (MEX) file.

INTRODUCTION

CAFE (Channel Access interFacE) [1, 2] is a modern, C++ client library that provides an intuitive and multifaceted user interface to the EPICS (Experimental Physics and Industrial Control System) [3] native C-based Channel Access (CA) Application Programming Interface (API) [4]. It allows for remote access to control data, encapsulated in Process Variables (PVs) residing in EPICS Input/Output Controllers (IOCs), while sheltering the user from the intricacies of programming with the native CA library. CAFE's conception arose from requirements foreseen by SwissFEL, Switzerland's X-ray Free-Electron Laser [5], coupled with the desire to avoid deprecated CA APIs propagating into a new project. Its development has since encompassed a renewed effort as requirements for application development for the forthcoming commissioning phase became apparent [6]. In particular, it was recognized that the effort afforded to a complete API, that further provided abstract layers for beam dynamics applications, could be readily incorporated into any C/C++ based scripting languages of choice. The main advantages to this approach are:

- The inherent simplicity and convenience of maintaining a single CA interface code. New CA functionalities from future EPICS 3 releases need only be integrated into a single base library.
- A uniform response to errors and exceptions that facilitates trace-backs.
- The CA class is well separated from the internals of the domain language meaning that bindings to other scripting and domain-specific libraries are vastly simplified.

An overview of the structure of the code is given, together with a discussion on design features that provide control over configurable components that govern the behaviour of interactions, and guarantee that the outcome of all method invocations are captured with integrity in every eventuality. Newly created bindings to Cython are then presented, and the improvement in performance yielded to Python developers is emphasized. An updated version of CAFE's MATLAB Executable (MEX) file [7], that exposes new functionalities to MATLAB [8] users, has also been made available.

CAFE C++ IMPLEMENTATION

The C++ interface to the EPICS Channel Access client library follows sound practices in CA programming [9] by placing careful attention to:

- Management of client-side CA connections.
- Memory optimization, particularly when connections are restored.
- Separation of data retrieval from its presentation.
- Strategies for converting between requested and native data types.
- Caching of pertinent data related to the channel and its state.
- Aggregation of requests for enhanced performance.
- Adaptive correction procedures, e.g. for network timeouts.

CAFE provides functionality for synchronous and asynchronous interactions for both single and groups of channels. All transactions report their data to a multi-index container provided by the Boost C++ libraries [10]. Here, the container takes ownership of instances of the "Conduit" object, each of which acts as the storage location for all data related to its associated PV. Multiple, distinct interfaces (or indices) provide convenient access to the object elements, allowing their data to be quickly retrieved or modified. The handle index (or object reference) has been configured with a unique key, and as such, acts at the definitive reference to the resource's data. To facilitate data access the underlying container is also indexed by PV, PV Alias, and the CA Channel Identifier. Callback functions have been implemented on all operations involving connection handlers, event handlers and access right handlers. Their invocation triggers their data to be written into the "Conduit" container object. In this way, the connection state of the channel, and all the channel's parameter values, are recorded with integrity. Memory to hold the channels data is allocated dynamically on first connection, and re-examined in the event of re-connection. Only one connection per channel is ever established, unless the channel is also member of a synchronous group (which itself

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

TINE STUDIO, MAKING LIFE EASY FOR ADMINISTRATORS, OPERATORS AND DEVELOPERS

P. Duval, M. Lomperski, DESY, Hamburg, Germany
J. Bobnar, Cosylab, Ljubljana, Slovenia

Abstract

A mature control system will provide central services such as alarm handling, archiving, location and naming, debugging, etc. along with development tools. It has become common to refer to the collection of these services as a 'studio'. Indeed Control System Studio (CSS) [1] strives to provide such services independent of the control system protocol. Such a 'one-size-fits-all' approach is likely, however, to focus on features and behavior of the most prominent control system protocol in use, providing a good fit there but perhaps offering only a rudimentary fit for 'other' control systems. TINE [2] is for instance supported by CSS but is much better served by making use of TINE Studio. We report here on the rich set of services and utilities comprising TINE Studio.

INTRODUCTION

What one ends up referring to as a control system *studio* is primarily all about *the user*. In a perfect world *control* of a machine could be *blind*, where input simply leads to output and there is no necessity of involving something as grandiose as a *studio*. In reality, this will never be the case. Namely, a user operating a machine will inevitably encounter problems which need to be quickly diagnosed and corrected. A user studying the machine will need to readily access all machine parameters. A user developing controls for a machine will also need to produce robust, functional and intuitive applications in a timely manner. A user administering a machine will need to be able to check the integrity of the controls and make modifications where necessary. We refer to the tools the user has at his disposal toward these ends as a *studio*, and where the control system is TINE, as TINE Studio.

The operator and machine physicist will make good use of TINE central services. Namely, the archive and alarm systems and their peripheral applications are major workhorses in machine control.

Likewise, developers will make use of server *wizards* and configuration editors, where front-end and middle layer servers are concerned, and rapid application development (RAD) components and GUI builders, where client applications are concerned.

Finally, control system administrators will profit from the many database management tools, report generators, and monitoring tools available.

In each case the utilities available strive to present an intuitive *cut-to-the-chase* application and take pains to avoid *popup pollution* and needless searching for results in arcane tree structures.

CENTRAL SERVICES

Archive System

An archive system can mean different things to different people. In the context of machine control, we will not refer to the data acquisition system necessary to process the myriads of data taken at an experiment. Instead we refer to the archiving of machine data versus time or versus event. The trend of relevant machine parameters can help diagnose operational problems when it is easy to browse and correlate among them. Likewise, investigating data and settings at a specific event, be it post-mortem, such as an RF trip, or otherwise is also invaluable.

TINE allows the central archiving of any machine parameter with optional storage criteria and does not modify, reduce, or expunge the data in any way after storage. A TINE server may also provide *local* archiving of any parameter.

The TINE archive viewer uses sparse (raster) data with points of interest for rapid browsing over very large time ranges. It also makes use of *optical zooming* whereby a new data request is made whenever a zoom operation over the time interval is made: the finer the zoom, the finer the raster. At any stage the number of points in the selected time range plotted versus the number of points available is easily seen. The archive viewer also decides on the optimal source (central or local) from which to retrieve the archive data.

An archived parameter can be of any available format, be it scalar, multi-channel array, waveform trace array, or video image. The primary archive view is, however, a trend, i.e. a value versus time. Consequently if an array element is selected for viewing then a trend of that element is displayed. It should be noted that the TINE archive viewer can also display the trend versus a system stamp (which might be a pulse or cycle number) instead of a time stamp. In any case, a secondary chart is able to display a snapshot of the entire array (multi-channel or waveform) at the selected (clicked on) time in the trend chart. Similarly, video images will display a frame count in the primary trend chart and the entire image in the secondary chart.

Figure 1 shows the trends of selected parameters in PETRA, including the horizontal tune, over a 24-hour period. The secondary plot shows horizontal tune at the time (green line) clicked on by the operator.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

APPLYING SOPHISTICATED ANALYTICS TO ACCELERATOR DATA AT BNL'S COLLIDER-ACCELERATOR COMPLEX: BRIDGING TO REPOSITORIES, TOOLS OF CHOICE, AND APPLICATIONS *

K.A. Brown†, P. Chitnis, T. D'Ottavio, J. Morris, S. Nemesure, S. Perez, D. Thomas
Collider-Accelerator Department., BNL, Upton, NY

Abstract

Analysis of accelerator data has traditionally been done using custom tools, either developed locally or at other laboratories. Much of the data analysis is done in real time when the data is being logged. However, sometimes users wish to apply improved algorithms, look for data correlations, or perform more sophisticated analysis. In recent years we have investigated the use of tools to bridge standard analysis systems, such as MATLAB®, R, or SciPy, to the controls data repositories. In this report we will discuss the tools used to extract data from the repositories, tools used to bridge the repositories to standard analysis systems, and directions we are considering for the future.

INTRODUCTION

Historically our data repositories have not been that large or complex. Over time these repositories have grown and the kinds of data have become more complex. Today it is not unusual to be working with multidimensional arrays or images, as well as simple scalar data. But more significantly, we store far more data than we analyze [1]. If we want to mine that data and look for particular patterns or do complex correlations, our legacy tools are inadequate.

The development of more sophisticated data collection and analysis systems started after 2000, once the Relativistic Heavy Ion Collider (RHIC), within the Collider-Accelerator Complex (C-AD) at BNL, became operational [2]. The new systems that were put in place standardized the data formats across all data sources and placed the data into well-defined filesystem hierarchies. General tools were built to access all data in these repositories.

The operation of RHIC and associated accelerator systems requires comprehensive data logging systems, collecting and storing measurements from many physical systems. Analysis of this data is a critical part of any troubleshooting process and involves the detection and study of composite behavior patterns. Since most of the tools in place to do this are custom built applications with limited functionality, bringing in other systems, such as MATLAB® or R, allows more sophisticated analytics to be applied to the data. However, the main issue isn't so much getting systems that can do sophisticated analytics, but easily linking such systems to the data that has been stored.

* Work performed under Contract Number DE-SC0012704 with the auspices of the US Department of Energy.
† kbrown@bnl.gov

Data Collection and Storage

One step towards making use of other analysis tools is to get the data already collected and stored to be available to any application. To do this we make use of HTTP protocol multithreaded data servers [3].

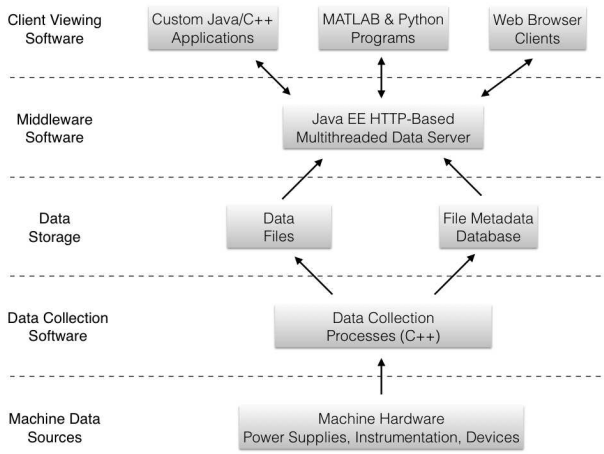


Figure 1: Three tier data layers for Accelerator Data.

Figure 1 shows how data gets collected and stored and then accessed using a data server. All the systems used at the *Machine Data Sources*, *Data Collection Software*, and *Data Storage* layers are legacy systems, for the most part, and serve the purpose of writing data into the filesystem. The *Data Server* provides a simple http-based interface to that data. We also have library level tools that provide this access and tools for obtaining live data.

Data Formats

The RHIC file storage and retrieval systems were designed to make use of the Self Describing Data Sets (SDDS) system developed at Argonne [4]. Using this file format a directory tree hierarchy was built, using a (standard relational) database to facilitate retrieval.

Today, the SDDS format remains the system used to store data, but that choice in format is now a constraint to new design choices. This is an issue in places where performance improvements need to be made. A lesson we seem to keep learning is that the opportunities of one era become the liabilities of a later era. Every custom design eventually becomes either a maintenance burden or a hurdle to progress [5].

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

USING THE VAADIN WEB FRAMEWORK FOR DEVELOPING RICH ACCELERATOR CONTROLS USER INTERFACES*

Wenge Fu#, Kevin Brown, Ted D'Ottavio, Seth Nemesure, Enrique Schuhmacher
Brookhaven National Laboratory, Upton, NY 11793, USA

Abstract

Applications used for Collider-Accelerator Controls at Brookhaven National Laboratory typically run as console level programs on a Linux operating system. One essential requirement for accelerator controls applications is the bidirectional synchronized IO data communication. Several web frameworks have made it possible to develop web based Accelerator Controls applications that provide all the features of console based user interface applications. Web based applications give users flexibility by providing an architecture independent domain for running applications. Security is established by restricting access to users within the local network. Additionally, the web framework provides the opportunity to develop mobile device applications that makes it convenient for users to access information anywhere and anytime. The Vaadin Java Web Framework is a tool kit being used to develop client side web interfaces. Vaadin provides Java developers a short learning curve overhead. Most Java Technologies, including JavaEE and third party packages work well within the Vaadin framework. This paper explores the feasibility of using the Vaadin web framework for developing UI applications for Collider-Accelerator controls at Brookhaven National Laboratory.

INTRODUCTION

"*Vaadin Framework is a Java web application development framework that is designed to make creation and maintenance of high quality web-based user interfaces easy*"[1]. First released in 2009, the Vaadin web application framework has been a fast growing API in terms of popularity among web developers for its rich functionality. The Vaadin framework has an advantage over other web development technologies because it uses the Java programming language which is more familiar to the application development community.

The Vaadin application framework provides two programming models: server side (Java) and client side. The client side framework is backed by the Google Web Toolkit (GWT). Program code is written in Java and resides on the server side. Server side program code helps make web applications more secure. Vaadin has a rich set of UI components. The server side and client side communicates via HTTP (or TCP when websockets are used) protocol and transfers data in JSON. The Vaadin framework supports all major web browsers without additional plugins[2] and works well with major IDEs. This makes the web application coding and debugging easier. In Vaadin, the look and feel of the web application is controlled by CSS themes. This makes web application GUI richer, and more configurable; Vaadin is best used for designing single page web applications which typically work like console level UI applications.

* Work supported by Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy.

fu@bnl.gov

VAADIN FOR ACCELERATOR CONTROL APPLICATIONS

Controls applications used in accelerator controls systems have many common characteristics, such as:

- They mostly require fast live bidirectional communications with many different systems such as hardware controllers, database servers, file servers, and other legend systems on different platforms (Unix, Linux, Windows etc.);
- Requires fast UI and interactive responsiveness.
- Rich UI for control data visualization for single or multiple GUIs.

These features can be relatively easy to implement with traditional languages such as C++ and Java. As Vaadin uses the Java programming language, server side java JDK (or JavaEE) APIs, third party APIs and jar packages can be used directly. This makes Vaadin a favorable choice when choosing a web application framework.

In a Vaadin web application, the server side programs (written by developers) and client side code (generated by Vaadin from server side code) have a common shared state, which helps enhance UI responsiveness and overall performance. For web based accelerator controls applications, the UI design and GUI layout are relatively easy to implement. It is critical for the client-server bidirectional communication layer to be effectively managed. Vaadin data push features make this kind of bi-directional communication easy to setup.

Vaadin push can be configured with Java annotations: @Push(PushMode, Transport) in program code or with a configuration file. There are three push modes:

- AUTOMATIC (default)
- MANUAL
- DISABLED

and 3 transport methods for communication in the request/response cycles:

- LONG_POLLING
- STREAMING
- WEBSOCKET

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

DEVELOPMENT OF iBEACON BASED EQUIPMENT INVENTORY SYSTEM AT STAR EXPERIMENT

J. Fujita, M. Cherney,
Department of Physics, Creighton University, Omaha, NE 68178, USA

Abstract

An inventory system using iBeacon technology has been developed. Using a specially written iOS app, makes the location of the equipment easier to a workers during the routine access to the experiment. The use of iBeacons and iOS devices allow us to distinguish one equipment rack from another very easily. Combined with 2D barcode, the use of iBeacons may provide better inventory management of the equipment for experiments.

INTRODUCTION

STAR Collaboration (Brookhaven National Laboratory, Upton, NY) composed of 56 institutions from 11 different countries. During an annual data-taking period, nearly a hundred different members in the collaboration will come to participate in taking data each year. During the data-taking period, it sometimes is necessary to perform a hardware reset. While the experts of each subsystem are typically present during the data-taking period, not every expert is available for emergencies, and the reset must be handled by the workers. Not all workers are familiar with the hardware, and they may not be aware of where the hardware equipment is located. Having a method to guide them to the hardware which needs to be reset is be critical.

To accomplish this, we have briefly looked into the possibility of using GIS. In ICALECPS 2005, Larrieu et al. [1] and in ICALEPCS 2007, Yamashita et al [2] presented the potential use of GIS for accelerator site management. Unfortunately, GIS is not always useful in three-dimensional space in close proximity such as a high energy physics experiment detector. The idea was quickly abandoned.

Recently, with the widespread use of mobile phone with Bluetooth Low Energy technology, utilizing iBeacons can be used to locate equipment in the close proximity. We have developed a special iOS app using iBeacons as well as a 2D barcode.

THE SYSTEM

For the iBeacons, we deployed several Raspberry Pi [*3*] computers equipped with USB Bluetooth 4.0 dongles.

SYSTEM OVERVIEW

iBeacon

iBeacon is the technology standard developed by Apple, which allows mobile apps to listen for signals from beacons in the physical world and react accordingly. Essentially, iBeacons behave as lighthouses. Using Bluetooth Low Energy technology, it sends out a small string of information to identify the beacon. This information is received by a specifically written app for the beacon to show the information as well as measure the distance from the beacon to the mobile phone.

iOS Application

In order to effectively use iBeacons, a special app that has the beacon information has to be written. In our case, the equipment rack information is associated with the beacon information. Once the worker enters the proximity of the equipment rack with an iBeacon, the app will notify the worker what equipment are available in that particular rack. This will let the worker know that he/she is standing in front of the correct set of equipment.

QR Code

While iBeacons will help the users to know if they are in the proximity of where they need to be, it will not help them to identify the equipment in question. To deal with that, QR codes are used. The information on the QR code is kept relatively short, allowing us to miniaturize the QR code size, which is necessary to fit it on a VME face plate. The built-in camera on an iOS device functions as a QR code scanner. Figure 1 shows an example use of the QR code.



Figure 1: An example of equipment with the QR code.

This implementation provided both flexibility and the cost effectiveness.

The iBeacons were placed throughout the equipment platform. They are set to transmit slightly different

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

ADOPTING AND ADAPTING CONTROL SYSTEM STUDIO AT DIAMOND LIGHT SOURCE

M. Furseman, N. Battam, T. Cobb, I. Gillingham, M. T. Heron, G. Knap, W. Rogers,
Diamond Light Source Ltd, Oxfordshire, UK

Abstract

Since commissioning, Diamond Light Source has used the Extensible Display Manager (EDM) to provide a GUI to its EPICS-based control system. As Linux moves away from X-Windows the future of EDM is uncertain, leading to the evaluation of Control System Studio (CS-Studio) as a replacement. Diamond has a user base accustomed to the interface provided by EDM and an infrastructure designed to launch the multiple windows associated with it. CS-Studio has been adapted to provide an interface that is similar to EDM's while keeping the new features of CS-Studio available. This will allow as simple as possible a transition to be made to using CS-Studio as Diamond's user interface to EPICS. It further opens up the possibility of integrating the control system user interface with those in the Eclipse based GDA and DAWN tools which are used for data acquisition and data analysis at Diamond.

INTRODUCTION

Diamond Light Source is a third generation light source, comprising of an injection chain of 100 MeV linac, 3 GeV booster ring and a 561.6 m 3 GeV storage ring [1]. There are currently a total of 33 photon beamlines which are either completed, in construction, or planned. All control system parameters are exposed via the Experimental Physics and Industrial Control System (EPICS) [2], and GUIs for these are realised with the Extensible Display Manager (EDM) [3], which is widely used in the accelerator community for monitoring and controlling live process variables (PVs).

EXISTING INFRASTRUCTURE

At Diamond Light Source we use EDM extensively across the Accelerator complex and Photon Beamlines, which allows us to provide a common interface to many core components such as vacuum systems and motors. Typically a user will access synoptic overviews of a beamline or the machine areas from the Diamond Launcher, an application that resembles a 'start menu' with Diamond specific content. From the synoptic screen a user will access more detailed information by clicking through a series of panels.

All the panels are stored on a read only file system which contains multiple versions of releases for different modules. EDM is started from a script invoked by the launcher which sets environmental variables to point at the correct versions of dependencies for any particular synoptic screen. A new instance of EDM is run for each synoptic allowing them to depend on different modules. In total around 7000 EDM screens are installed at Diamond with many of these being auto-generated.

Not all displays in Diamond are thin clients that interact only with EPICS PVs; some require additional processing to provide the user with interactive information or need to calculate values based on a physics algorithm. It is not plausible to do these processes in EDM and is sometimes difficult to achieve them in an IOC. These are typically implemented in PyQT and interact with EPICS using the Cothread Python library [4]. Developing these GUIs is time consuming and can provide an inconsistent feel to the overall operator interface.

MOTIVATION FOR MOVING TO CONTROL SYSTEM STUDIO

While Diamond has been using EDM since it began operation, the long term prospects for the application are uncertain. As Linux distributions move away from X to Wayland, EDM will lose support for Motif, the widget toolkit it is built on. While it is plausible that EDM could be written to use a more modern toolkit, this transition provides the opportunity to look for a new GUI tool for creating operator interfaces that are based on more modern technologies and provide more features.

Control System Studio [5] (CS-Studio) is written in Java and based on the Eclipse Rich Client Platform (RCP). It provides a modular plugin architecture that allows extensions to be easily contributed in the form of plugins. This architecture has given rise to a number of features that can be bundled with the application to tie together a suite of control system tools into a single interface; an example of one of these plugins is the Data Browser, which can be used simultaneously as a replacement for the 'Striptool' and the 'Archive Viewer'. Eclipse RCP is capable of running on multiple platforms including Linux and Windows. While our control interfaces run on Linux workstations, this does remove a barrier in providing control system information to office users with Windows workstations.

CS-Studio is the choice for many new sites [6–8]. This increases confidence that the application will benefit from community maintenance in the future. Indeed, there has been an average of 7.82 commits a day to the master branch since 2007 [9].

Yet another benefit of moving to CS-Studio is that an automated conversion framework already exists for converting EDM layout files (EDL) to CS-Studio's XML based layout files (OPI). With so many EDL files automated conversion is a necessity if Diamond is to transition its entire GUI infrastructure in a reasonable time scale. The conversion framework is also easily extensible allowing developers to add unimplemented features.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

TOOLS AND PROCEDURES FOR HIGH QUALITY TECHNICAL INFRASTRUCTURE MONITORING REFERENCE DATA AT CERN

R. Martini, M. Braeger, J. Salmon, A. Suwalska CERN, Geneva, Switzerland

Abstract

The monitoring of the technical infrastructure at CERN relies on the quality of the definition of numerous and heterogeneous data sources. In 2006, we introduced the MoDESTI procedure for the Technical Infrastructure Monitoring (TIM) system to promote data quality. The first step in the data integration process is the standardisation of the declaration of the various data points whether these are alarms, equipment statuses or analogue measurement values. Users declare their data points and can follow their requests, monitoring personnel ensure the infrastructure is adapted to the new data, and control room operators check that the data points are defined in a consistent and intelligible way. Furthermore, rigorous validations are carried out on input data to ensure correctness as well as optimal integration with other computer systems at CERN (maintenance management, geographical viewing tools etc.). We are now redesigning the MoDESTI procedure in order to provide an intuitive and streamlined Web based tool for managing data definition, as well as reducing the time between data point integration requests and implementation. Additionally, we are introducing a Class-Device-Property data definition model, a standard in the CERN accelerator sector, for a more flexible use of the TIM data points.

INTRODUCTION

Computerised monitoring of the technical infrastructure at CERN goes back to the early 1980s when a DEC DPD11 was used to display around 2000 alarms onto 2 screens in a dedicated control room. When this system was migrated to a distributed solution based on Unix PCs the decision was taken to hold monitoring definition data on an external database from which the monitoring system would be configured.

This strategy was of great benefit as the monitoring system evolved first to the Technical Data Server – (TDS) in the mid 1990s and finally to the Technical Infrastructure Monitoring (TIM) system [1]. Throughout this time the number of data points grew to reach around 100K data points today, and there is a constant demand to incorporate new monitoring from different services installing new equipment and upgrading existing installations.

For all these monitoring systems the issue of ensuring that data were correctly defined was paramount in determining satisfactory performance. The specification of the monitoring data is the responsibility of the different services whose equipment is being monitored, but the data must be validated by the operators who monitor the alarm screens as well as those responsible for the monitoring system itself.

To ensure that the declared data are correct, complete and follow established standards, the Monitoring Data Entry System for Technical Instructure (MoDESTI) procedure [2] was devised and deployed shortly after the implementation of TIM in 2005.

MoDESTI has evolved over the years to handle the many different scenarios relating to the maintenance of the data that defines TIM; however it is cumbersome to use and difficult to maintain when changes in the monitoring system need to be covered. For this reason a complete reworking of the tool was proposed in 2014 and is now in development.

MODESTI PRINCIPLES

Templates Monitoring data includes descriptive information such as location, equipment concerned, person responsible, information in dealing with an alarm (priority, causes, consequences, actions to take), applications that use the data point (synoptic views, logging, external systems), and configuration data that indicates which acquisition units handle the data points as well as how they are communicated to and from TIM. In this way the reference database contains an overview of all that is monitored by TIM and can be consulted by any authorised user.

Changes will occur during the lifetime of a data point, either through corrections to inaccuracies or due to changes in the way the data point is monitored. Furthermore data points are not monitored indefinitely, when an installation is dismantled, the related data points must be removed from the monitoring system to avoid clutter and possible confusion. These events are also managed by MoDESTI so that the monitoring systems always keep up with the latest operational changes.

Since the aim of MoDESTI was to promote data quality in monitoring, the tool was adapted for the launch of the CERN Safety Alarm Monitoring (CSAM) system. This is a separate system from TIM but CSAM is configured with data defined through the MoDESTI procedure held on the reference database. More recently, the capacity of MoDESTI in enforcing standard definitions and rules relating to alarms was harnessed by CERN's WinCC application for alarms sent to the CCC.

In order to declare data points for monitoring by TIM users must create a data integration request and submit their data in a standard way. Depending on the nature of the data points being declared, specific actions must be taken by the different agents involved in the monitoring process. Alarms must be validated by CCC operators to verify that they are correct, complete and comprehensible

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

ADVANCED MATLAB GUI DEVELOPMENT WITH THE DataGUI LIBRARY

Sascha Meykopff, DESY, Hamburg, Germany

Abstract

On the DESY campus Matlab is a widely used tool for creating complex user interfaces. Although the on-board GUI tools are easy to use and provide quick results, the generated low-level code lacks uniformity and advanced features like automatic verification and conversion of input and output data. These limitations are overcome by the newly developed DataGUI library. The library is based on the model-view-controller software pattern and supports enhanced data handling, undocumented Matlab GUI elements, and configurable resizing of the user interface. An outlook on features of the upcoming release is also presented.

INTRODUCTION

At the European XFEL and the FLASH facility, both located on the DESY campus, most of the operation handling is based on the client-server model with JDDD as the front-end client and DOOCS servers as the back end software [1]. If an operation task needs a more complex graphical user interface (GUI) it's recommended to use Matlab. The GUIDE tool is an interactive method to create a GUI with Matlab. The use of GUIDE provides quick results but extensive code work is necessary to handle the interaction between the GUI elements and the program logic. The effort to implement a GUI complete programmatically without the help of GUIDE is small. The DataGUI library was developed to reduce the coding amount, and to improve the stability of Matlab applications.

ACCESS AND VERIFY PROPERTIES

Matlab GUI elements are designed as objects. The GUI element behavior is defined by object properties. To access element properties the developer has to call a get or set function. Starting with Matlab version R2014b one can use the dot notation to query or set properties [2]. The property names and type format depend on the GUI element type and Matlab version. To write an application which works on different Matlab releases these circumstances must be considered.

```
function edit_field_callback(handle, event)
   edit_input = get(handle, 'STRING');
   edit_input = str2double(edit_input);
   if ~isnan(edit_input)
      display(['input is ' num2str(edit_input)])
   else
      display('input invalid')
   end
```

Figure 1: How to verify and convert text field input in a Matlab callback function.

In a good software design every input variable should be verified. To neglect parameter verification results in unstable software. Currently a Matlab developer is not encouraged to verify the GUI input. Figure 1 shows an example how to verify and convert an edit text field for integer input. This code has to be implemented for every input element.

SHARE DATA AMONG CALLBACKS

The interaction between GUI elements and user code is done by callback functions. If one clicks on a GUI button, Matlab calls an assigned function. Most of the Matlab callbacks carry two input and no output parameters. The first parameter is the handle of the triggered GUI element. The second parameter is defined as 'eventdata' without any value. The developer can add additional input parameters. No return parameters are defined in Matlab callbacks. If the code needs to share data among callbacks the function parameters are not suitable. The Matlab documentation describes four different solutions [3]. These solutions have in common to use a function to get and store the shared data. These functions needs to be called at the start and the end of every callback. This design causes a lot of code duplication and potential errors. A lot of real world software use global variables to store data among callbacks. This is one of the worst solutions because global variables yields in a bad code design.

MODEL-VIEW-CONTROLLER

The Model-view-controller (MVC) is a software pattern which is followed by the most modern GUI toolkits. The pattern is an abstract idea of separating the code into three parts. One part describes and handles the data model of an application. The second one creates the view of the data model. Multiple views of the same data model are allowed for example to view a data table and a plot. The last part covers the program logic. This part updates the data model and modifies the different views. The MVC pattern gives a rough software structure and helps the developer to create a good software design. The software design usually suffers because developing a Matlab GUI is time-consuming. The DataGUI library guides the developer to separate his code which results in a better design.

GLOBAL DATA DESCRIPTION

Matlab software who uses the DataGUI library has unified callbacks. Every callback has a 'data' variable as input and output. The type of the 'data' variable is a structure. One can add, remove, and store fields inside the structure. The storage among callbacks is handled by the

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

A STRUCTURED APPROACH TO CONTROL SYSTEM GUI DESIGN FOR THE SOLARIS LIGHT SOURCE

V. Juvan, I. Dolinšek, T. Humar, M.Pavleski, Cosylab, Ljubljana, Slovenia
P. P. Goryl, Solaris, Krakow, Poland

Abstract

In the framework of delivering control system services to the Solaris synchrotron light source, Krakow, Poland, Cosylab realized a comprehensive set of controls GUIs, using a structured approach. The goals of using this architecture are threefold. The first is to achieve reliable, predictable and consistent behaviour of the controls software. The second is that it is easy to deploy and maintain through scripting. The third is that it is future-proof by providing extensibility, using dedicated templates. The system is based on a configuration database, populated with devices, device specifics and device groups (clusters of devices performing specific operations). The GUIs are dynamically generated from this configuration. For the synoptic views, TANGO-standard JDraw and its configuration are integrated into the framework. Existing GUIs, written in PyTango can be easily adapted to function as part of this system. The compelling user benefits are high usability and life-time management through controlled upgrade and extension. For new big physics projects this GUI control program offers a customizable solution for any TANGO based control system.

CHALLANGES

When realizing the controls GUIs, we were presented with a series of challenges that we had to overcome.

Firstly, we had to develop a full machine control and state overview program. It was designed as a single entry point program for performing all operations, from every day operator to expert use. Moreover, a broad set of controls GUIs was required for specific operations.

Multiple instances of the control program were foreseen, each intended for a separate control system at the facility.

During the development, we expected to adjust the functionality and features to user feedback, which would be acquired in steps.

Existing tools were not sufficient since they did not provide the desired functionality. With a limited budget, we undertook a development of the Solaris Synchrotron Control Program.

GOALS

The Solaris Synchrotron Control Program had to meet a variety of goals.

We wanted to achieve a predictable and consistent behaviour of the control room software. The control program had to provide a required set of functionalities and features for operation. A transparent and convenient use of the software was requested.

A very important aspect of the control program was that it had to be configuration driven. The goal was to maintain one source of information from which the control program would gather all required information. We would evade hardcoding at all costs. Moreover, we made use of the dynamic GUI generation and generic panels to enable the control program to be adjustable by only modifying its configuration source.

A significant goal was to guarantee easy extensibility. This was achieved by providing dedicated templates and support for integrating external applications in the control program itself.

Last but not least, we wanted to provide easy deployment and maintenance through scripting.

SOFTWARE ELEMENTS

The final deliverable consists of four components, namely:

- Configuration database:
  - Device configuration
  - Device group configuration.
- Control Program.
- Custom GUI library
- GUIrunner

Configuration

The configuration was split into two parts, the device configuration and the device group configuration. The first holds all the information regarding the Tango devices that were to be managed by the control program. The latter holds the information in respect to the groups of devices, to which they may belong to. It heavily relates to the custom GUI library.

The device information incorporates the following parameters:

- Element name
- Type
- L[m], S[m], X[m], Y[m], Z[m],
- Section
- Subsystem
- Managed in CS
- Device Server name
- Device Server instance
- Device class
- Full TANGO device name
- Alias
- Triggered by TTL
- Custom GUI
- Aggregate GUI

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

GUI STYLE GUIDE FOR CONTROL SYSTEM APPLICATIONS AT ESS

Frank Amand, Miroslav Pavleski, Mark Plesko, Cosylab, Ljubljana, Slovenia
Leandro Fernandez, ESS, Lund, Sweden

Abstract

To help developers create consistent-looking control system application GUIs, the European Spallation Source Integrated Control Systems group asked Cosylab to develop a Style Guide (SG) document.

Its purpose is to avoid that GUIs needlessly diverge and make the end-result of all screens combined look harmonious, even if GUIs have been developed over several years by many contributors.

Also it will speed up development, by letting developers start from design patterns, rather than starting "from a blank page".

The document defines a set of basic panel sizes, containing a 960px-style grid for consistent organization of content. It also defines a color scheme and font usage, in-line with the overall ESS corporate communications manual, with the addition of signal colors.

In addition it shows example screens to serve as GUI design patterns for typical screen types such as engineering screens, control applications and synoptic screens.

It concludes by setting rules and recommendations for the usage of automation symbols and display of engineering and physical units.

The document is further complemented by a separate document with Usability Guidelines for Human-Machine interfaces.

CONTROL SCREENS MATCHING ESS BRAND

The GUI colors and fonts that the SG defines were chosen to match he ESS color scheme as defined in the ESS corporate communications manual.

The ESS corporate style has a blue hue as the main color. This is a very suitable color for styling GUI elements. The choice to be in line with the ESS corporate style can enable the control GUIs to look nice and consistent by themselves.

An additional benefit is that this way the GUIs will also be harmonious when used in conjunction with web-based applications that are made in style of (use cascading stylesheet templates similir to) the ESS website itself.

DEFAULT PANEL SIZES AND A GRID

A grid is a proven technique for creating consistent looking GUIs e.g. in web-design.

The style guide defines a "960 pixel grid" in 12 colums. From this basis, 2, 3, 4 and 6 column layouts can be derived and combined.

The style guide also provides patterns for smaller screens/dialogs. For full screen applications it recommends combining smaller grid-based elements.

SIGNALS STAND OUT FROM BACKGROUND

Alarms and signal changes should be very noticable. Their full saturated colors clearly stand out from the *by-design* subdued backgrounds.

The included design patterns illustrate how to use moderate-sized LEDs and other indicators. This keeps a balance of bright colored signals versus the more sober backgrounds and it maintains the overall easthetics. (A Control GUI should by no means a Christmas Tree)

TITILLIUM + OPEN SANS FONTS

The SG defines Titillium [1] for titles (but for titles only). This Open Source font is the ESS corporate font and has nice distinctive features.

But to maximise readability for the body of the GUIs, also at smaller sizes, the more standard-styled Open Sans is chosen. This Google font is a modern, humanist sans-serif. It is designed for good readability at various font sizes. It is NOT just another Arial. It has more font weights (10 in total) then the old normal + italic + bold + bold-italic set that OS standard true type fonts used to provide.

DESIGN PATTERNS AND OTHER REUSE

Various design patterns reuse best practices from other projects on the design of

- Control screens: task oriented screens, either for monitoring and control of a single device or a contral application employing a combination of physical equipment / devices.
- Diagnostics screens, such as a beam current monitor.
- Synoptic views, showing machine overview in a graphical way.

The style guide also defines the way automation symbols are used. In this project we decided to reuase the efforts done on incorporating a symbols database in the EPICS CS Studio tool by the ITER CODAC initiative.

Some effort was spent on making recommendations for implementing efficient and user friendly machine browsing and naviation. This is often a painpoint in control systems of large machines.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

ALICE MONITORING IN 3-D

O. Pinazza, INFN Sezione di Bologna, Bologna, Italy; A. Augustinus, P. M. Bond, P. Chochula, L. M. Lechman, CERN, Geneva, Switzerland; A. N. Kurepin, INR RAS – Institute for Nuclear Research of the Russian Academy of Sciences, Moscow, Russia;
Jeremi Niedziela, Warsaw University of Technology, Poland, and CERN

Abstract

The ALICE experiment is a complex hardware and software device, monitored and operated with a control system based on WinCC OA. ALICE is composed of 19 detectors and installed in a cavern along the LHC at CERN; each detector is a logical set of modular elements, assembled in a hierarchical model called Finite State Machine. A 3-D model of the ALICE detector has been realized, where all elements of the FSM are represented in their relative location, giving an immediate overview of the status of the detector. For its simplicity, it can be a useful tool for the training of operators. The development is done using WinCC OA integrated with the JCOP fw3DViewer, based on the AliRoot geometry settings. Extraction and conversion of geometry data from AliRoot requires the usage of conversion libraries, which are currently being implemented. A preliminary version of ALICE 3-D is now deployed on the operator panel in the ALICE Run Control Centre. In the next future, the 3-D panel will be available on a big touch screen in the ALICE Visits Centre, providing visitors with the unique experience of navigating the experiment from both inside and out.

INTRODUCTION

With a length of about 26 m, width and height of 16 m, ALICE [1] is one of the biggest and most complex experimental apparatus in the world. It's installed at CERN in Geneva, at access Point2 of the Large Hadron Collider, in an underground hall at a depth of 52 m.

The experimental apparatus is extremely complex, being composed of 19 sub-detectors, each with its own specific technology choices. Monitoring, operating and managing ALICE is accomplished in collaboration by more than 1000 physicists and engineers, from 105 Institutes in 30 different countries.

The ALICE Detector Control System (DCS) [2] is a composite hardware and software structure, connecting and coordinating controls and operations on the sub-detectors and the infrastructure. All software controls are developed and integrated through the SCADA software framework WinCC OA [3], selected by CERN as a standard for all experiment control systems.

Monitoring and controlling ALICE through the DCS is performed by an on-site shifter on a 24/7 basis, whose main duties are to guarantee the safety and the integrity of the detectors and the infrastructure, and smooth running of the experiment. Several on-call experts collaborate with the shifter for the most delicate operations.

The interfaces available for the operators are mainly table based, and all states and actions are encoded using a Finite State Machine (FSM) concept [4]. In this way, complexity is hidden and operators are able to evaluate the status in real time, based on simple colours and keywords.

3-D GRAPHICS FOR WINCC-OA BASED CONTROL SYSTEMS

In spite of the evident simplicity of the tabular panels, a more realistic view of the system with its geometrical characteristics would allow a better understanding and a more faithful evaluation of the overall status of the experiment. For this reason, we decided to develop a new panel and display the sub-detectors and their actual FSM state in an intuitive way (see Figure 1).

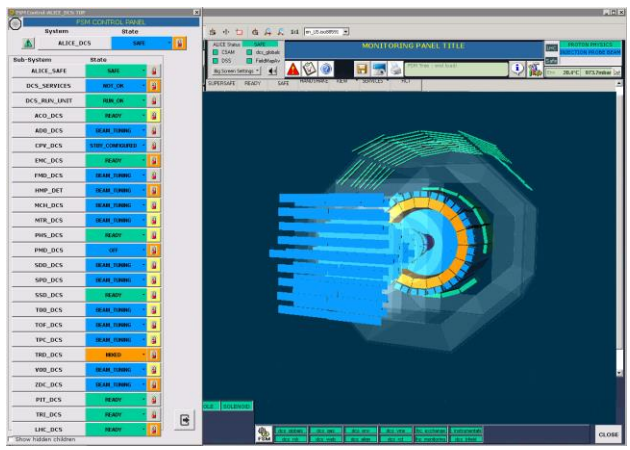


Figure 1: Tabular and 3-D representation of ALICE in one of the operator panels.

All DCS operator panels are integrated in WinCC OA. On top of WinCC, CERN has delivered a set of software tools called JCOP Framework, where common solutions and developments typical to HEP needs are made available. Among these tools, a 3-D viewer widget is available as an extension of the WinCC OA User Interface.

The fw3DViewer Framework tool [5] is based on the GEANT logic [6], and permits building several different shapes using the standard scripting language available in WinCC. The shapes can be simple "Boxes", "Spheres", "Cylinders", "Tubes", "Cones", up to more general "Trapezoids" and "Polyhedra". Each shape is represented by a typical set of geometrical data: a sphere is defined by the position of its centre and the radius; drawing a box

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

UNIFYING ALL TANGO CONTROL SERVICES IN A CUSTOMIZABLE GRAPHICAL USER INTERFACE

S. Rubio-Manrique, G. Cuní, D. Fernández-Carreiras,
C. Pascual-Izarra, D.Roldán, ALBA Synchrotron, Cerdanyola del Vallés, Spain
E. Al-Dmour, Max-IV, Lund, Sweden

Abstract

TANGO is a distributed Control System with an active community of developers. The community features multiple services like Archiving or Alarms with an heterogeneous mix of technologies and look-and-feels that must be integrated in the final user workflow. The Viewer and Commander Control Application (VACCA) was developed on top of Taurus to provide TANGO with the user experience of a commercial SCADA, keeping the advantages of open source. The Taurus GUI application enables scientists to design their own live applications using drag-and-drop from the widget catalog. The VACCA User Interface provides a template mechanism for synoptic-driven applications and extends the widget catalog to interact with all the components of the control system (Alarms, Archiving, Databases, Hosts Administration). The elements of VACCA are described in this paper, as well as its mechanisms to encapsulate all services in a GUI for an specific subsystem (e.g. Vacuum).

INTRODUCTION

Tango and Taurus

ALBA[1] Synchrotron is a third generation Synchrotron lightsource in Barcelona, Europe, providing synchrotron light since 2012 to users in its 7 beamlines, with 2 more under construction. ALBA institute has been an active member of the Tango Collaboration[2][3] since the very beginning of its design and construction phase. Tango is an open source object-oriented control system, done in collaboration between ESRF and a growing community of institutes and companies developing new device servers and tools using either C++, Python or Java.

Our Human Machine Interfaces to the Tango Control System are developed using Taurus[4][5], a framework for creating both GUIs and command-line tools to interact with scientific[6] and industrial control systems and related data sources. Originally developed in-house at ALBA, Taurus opened its development to the members of the Tango community, becoming popular among the new members of the Tango Collaboration. Some of the causes of this success are the technologies involved: Qt, Python[5] and its integration with SciPy[6], the popular stack of scientific libraries for python.

Tango as SCADA

The concept of a Supervisory Control and Data Acquisition system (SCADA) is largely used in both industrial and scientific worlds to define a control system that remotely controls large installations, using multiple communication channels and providing an application to control large processes distributed on many remote equipments or stations

Tango fits perfectly within SCADA definition, providing the communication channels and software tools needed to manage large particle accelerators and other facilities, along with additional control services such as Archiving, Alarms and User Access. These services are managed by a collection of applications developed either by the core team or by other members of the community: Jive for database configuration, Astor for control host management, JDraw for Synoptics, Mambo for Archiving, ATK/Taurus/QTango for GUI development, Sardana for experiment control as well as several web toolkits and multiple alternatives for Alarm systems[7] like PANIC[8].

Those applications provide a rich functionality and a full-featured control system, but also a diverse collection of look & feels and workflows that may be inconsistent and interfere with user interaction (Fig.1). This paper will address this issue providing a Taurus-based solution.

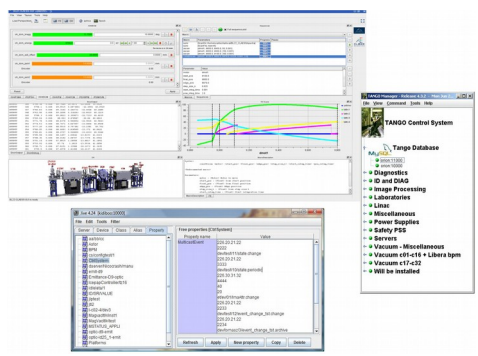


Figure 1: Tango look & feels: Taurus, Jive, Astor

State of the Art, Taurus and CSS

GUI consistency issues are not unique to Tango or to open source control systems, as they may apply to any collaborative project on which each member needs to tune the generic tools to its own context.

At ALBA, these problems were mostly solved once Taurus[5] became the default GUI framework for all our applications, thus creating the opportunity to unify the user interaction with the whole control system. Taurus does not enforce control-system specific conventions but is open instead to include multiple schemas and allows

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

A HTML5 WEB INTERFACE FOR JAVA DOOCS DATA DISPLAY

E. Sombrowski, R. Kammering, K. Rehlich, DESY Hamburg, Germany

Abstract

JAVA DOOCS Data Display (JDDD) [1] is the standard tool for developing control system panels for the FLASH facility and European XFEL. The panels are mainly started on DESY campus. For remote monitoring and expert assistance a secure, fast and light-weight access method is required. One possible solution is using HTML5 as transport protocol, because it is available on many common platforms including mobile ones.

For this reason an HTML5 version of JDDD, running in a Tomcat application server, was developed. WebSocket technology is used to transfer the panel image to the browser. In the other direction, mouse events are sent back from the browser to the Tomcat server. Now thousands of existing JDDD panels can be accessed from remote using standard web technology. No special browser plugins are required.

This article discusses the general issues of the web-based interaction with the control system such as security, usability, network traffic and scalability, and presents the WebSocket approach.

INTRODUCTION

JDDD is a common tool for designing and running control system panels at DESY [2,3,4,5,6]. Around 3000 control panels are started each day from DESY offices and in the control room.

Experience shows that many of these panels are needed by experts for remote assistance. The most comfortable access way is using a web interface, which is available on any PC and mobile device all over the world without the necessity of installing any additional software. The requirement for this interface is that no special web version of the panels should be needed to avoid double work. The communication should be fast and light-weight and has to work also in limited bandwidth environments.

A common technology for modern responsive web communication is AJAX (Asynchronous JavaScript and XML). Using AJAX, the web browser polls the server for data on each update. For the transfer of big data packages at high frequency – as it is needed for the JDDD web interface - the creation of a new HTTP connection every time would be a bottleneck. A better solution is a persistent connection like WebSockets.

WEBSOCKET TECHNOLOGY

WebSockets provide a protocol between client and server, which runs over a persistent TCP connection. A visual representation of such a communication is displayed in Fig. 1:

Handshake

Any network communication that uses the WebSocket protocol starts with an opening handshake. The client web browser sends a request to the server to upgrade the HTTP to WebSocket protocol and if the server supports WebSockets, it sends back a response saying: Ok I am upgrading your HTTP to a single TCP socket connection.

Bidirectional Messages

Through this open connection, bi-directional, full-duplex messages can be exchanged at a high frequency (simultaneously or back and forth). In the case of a JDDD WebSocket connection, panel images, which are created on the server side, are sent to the client and mouse events, which are produced in the browser window, are sent back to the server.

Closing the Channel

If the client or server closes the connection, the panel image creation is stopped on the server and all monitors to control systems values are removed.

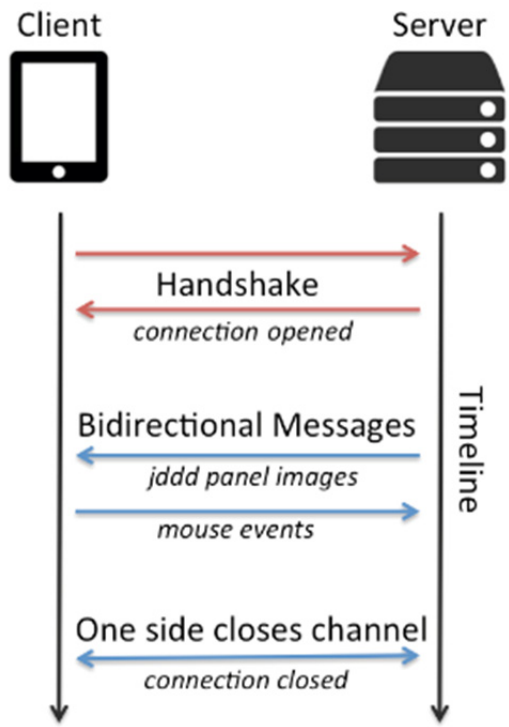


Figure 1: Using WebSockets for client / server communication.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

TIME TRAVEL MADE POSSIBLE AT FERMI BY THE TIME-MACHINE APPLICATION

Giacomo Strangolino, Marco Lonza, Lorenzo Pivetta, Elettra, Trieste, Italy

Abstract

The TANGO archiving system HDB++ continuously stores data over time into the historical database. The new time-machine application, a specialization of the extensively used save/restore framework, allows bringing back sets of control system variables to their values at a precise date and time in the past. Given the desired time stamp t0 and a set of TANGO attributes, the values recorded at the most recent date and time preceding or equalling t0 are fetched from the historical database. The user can examine the list of variables with their values before performing a full or partial restoration of the set. The time-machine seamlessly integrates with the well known save/restore application, sharing many of its characteristics and functionalities, such as the matrix-based subset selection, the live difference view and the simple and effective user interface.

TERMINOLOGY

Context

The context is a set of TANGO attributes gathered together under a name; the context structure is exactly the same as the one beneath the classical save/restore database.

Snapshot

The snapshot is always associated to a specific context and stores the values of the attributes at the very moment the snapshot itself was taken. The relation between a context and the snapshot is represented by the context ID and a snapshot has a unique snapshot ID. Each context can have multiple snapshots taken at different instants, while each snapshot is linked to a unique context. Of course, there is no constraint in the definition of the contexts: distinct ones can have one or more attributes in common.

WORKING PRINCIPLE

The classic save/restore [1] control room application was designed to save a snapshot of a set of TANGO attributes on a database at a given point in time. The snapshots can be later revised and restored on the field. Only previously saved sets of values, identified by an ID, a comment, the date and time the snapshot was taken, can be put back on the equipments. The time machine takes a different approach. Relying on an event-based TANGO historical database archiver, named HDB++ [2], it is able to restore a set of attributes to their values at any date and time in the past. The HDB++ constantly saves TANGO attributes into the database according to a given threshold. Taking advantage of the event-based archiver and the specific extraction library, it is always possible to retrieve the value of a quantity at an arbitrary point in the past. Historical database snapshots can thus be created on the fly at a given date and time.

THE HISTORICAL DATABASE LIBRARY

The HdbExtractor++ framework allows fetching data from a historical database in a simple object oriented fashion. The library currently supports the HDB (the original version of the TANGO archiver) and HDB++ MySql database schemas. Scalar and vector data can be fetched from the database, given a start and stop date and time. A separate module provides a Qt interface to the extraction library and also a dedicated GUI, aimed at displaying scalar and vector data over time. Extraction in the Qt module operates in a separate thread. Vectors are displayed over time by a surface plot, based on the MathGL framework. Errors are notified by NULL values. They typically reflect a read error from the device at the moment of the save process. If the data of a given source is not recorded in a given period, it is possible to fetch values older than the start date. Moreover, data sets pertaining to different curves can be aligned in time, so that a point on the curve is associated to each time stamp in the x axis. Even if the HdbExtractor++ library is tailored to deal with TANGO data (as far as data type, write mode property and data format are concerned), it is not dependent from TANGO specific types and libraries. The Qt module is multi threaded and can be easily integrated in graphical user interfaces [1].

THE GRAPHICAL USER INTERFACE

The Time-Machine graphical user interface is written in Qt and is made up of two sections: the classical save/restore and the new time machine mode. When the application is launched, the classical save/restore mode is selected by default (see Fig. 1). A radio button allows to switch to the time machine view (see Fig. 2). The design of the user interface unifies the approach to the selection of the contexts and the restore process in both modalities. The top left view is a list of the available contexts. Note that the save/restore contexts and time machine ones are defined in distinct databases, so in principle they don't overlap. The right view displays the TANGO attribute names and properties. The information is updated with the read and write values of each quantity as soon as a

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

Karabo-GUI: THE MULTI-PURPOSE GRAPHICAL FRONT-END FOR THE Karabo FRAMEWORK

Burkhard Heisen, Martin Teichmann, Kerstin Weger, John Wiggins,
European XFEL, Hamburg, Germany

Abstract

The Karabo GUI is a generic graphical user interface (GUI) which is currently developed at the European XFEL GmbH. It allows the complete management of the Karabo distributed control and data acquisition system. Remote applications (devices) can be instantiated, operated and terminated. Devices are listed in a live navigation view and from the self-description inherent to every device a default configuration panel is generated. The user may combine interrelated components into one project. Such a project includes persisted device configurations, custom control panels and macros. Expert panels can be built by intermixing static graphical elements with dynamic widgets connected to parameters of the distributed system. The same panel can also be used to graphically configure and execute data analysis workflows. Other features include an embedded IPython scripting console, logging, notification and alarm handling. The GUI is user-centric and will restrict display or editing capability according to the user's role and the current device state. The GUI is based on PyQt technology and acts as a thin network client to a central Karabo GUI-Server.

THE KARABO DISTRIBUTED CONTROL AND DATA ACQUISITION SYSTEM

Karabo is the control and data acquisition system which will be used on all beamlines of the European XFEL to control the equipment, acquire data and process it [1]. It is a centralized system where all communication is done via a central broker, except for high-bandwidth data streams which are transported via dedicated point-to-point connections.

A Karabo system is a collection of *devices*. Those devices can serve many purposes, they might be a driver for a particular hardware, a composite device that controls several other devices if some coordination between different hardware is necessary, a data processing device that may be just one of hundreds of equal ones in a server farm processing data, or a device saving raw or processed data to disk, not to mention many service devices which keep Karabo running.

There are many different approaches of a GUI for a control system. Some use an already existing development enviroment and extend them to the needs of developing GUIs (e.g. GDA [2], CSS [3]). Others developed stand-alone graphical editors (e.g. JDDD [4], Taurus [5]). Karabo has one fully integrated GUI, which is not only a graphical editor but can be used for all control and data acquisition tasks.

PROJECTS

A typical user of Karabo is not interested in the entirety of the system, rather a specific task to work on. Those tasks are often overlapping, as an example a vacuum technician may want to interact with the same components of an apparatus a scientist works on.

All information necessary for a task can be bundled into a *project*. These are all the devices needed and the configuration with which they should be started, or into which they should be reconfigured to perform the desired tasks. Graphical visualizations of tasks can be added as *scenes* to the project. They are used to show and edit the configuration of devices and the connections between them. Repetitive tasks reoccuring for users of a project may be programmed as macros. If taking some data is the purpose of a project, the data to be taken can be stored in the project as *monitors*. The project itself, however, does not contain the data, which is written to disk independently.

The projects are generally stored on a central server. This way everything needed for a task is persisted, and enables users to use their projects on different computers, and makes the administration of Karabo installations easier, as projects can be archived much simpler.

The project is the core concept of the Karabo GUI, and most of the rest of this paper is a description of its components. Figure 1 shows a screenshot of a running Karabo GUI with all components.

LIVE NAVIGATION AND CONFIGURATION

A Karabo device is run by a Karabo *server*, which is software running on a computer. The device code to be run is installed as plugins into those servers. The installation and running of those servers is beyond the scope of Karabo, but it is typically done with common server administration software.

In the GUI, such an installation is shown as a tree, for each computer the available servers are shown, and the available plugins for each server. A user can thus see the entire Karbo installation. From the GUI, the devices can be instantiated from their *device class*, the code to be run for a device.

Before a device is even instantiated, its parameters are already known to the device server and communicated to the GUI. Thus a user can configure a device interactively by filling out a form which is automatically generated from the description of the device class. Those configurations can be stored in the project. As the full live navigation tree of a Karabo installation can grow large and confusing, those stored configuration give a good overview of the devices needed in a particular project.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

VISUALIZATION OF INTERLOCKS WITH EPICS DATABASE AND EDM EMBEDDED WINDOWS

Evgeniy Tikhomolov, TRIUMF, Vancouver, Canada

Abstract

The control system for TRIUMF's upgraded secondary beam line M20 was implemented by using a PLC and one of many EPICS IOCs running on a multi-core Dell server. Running the IOC on a powerful machine rather than on a small dedicated computer has a number of advantages such as fast code execution and the availability of a large amount of memory. A large EPICS database can be loaded into the IOC and used for visualization of the interlocks implemented in the PLC. The information about interlock status registers, text messages, and the names of control and interlock panels are entered into a relational database by using a web browser. Top-level EPICS schematics are generated from the relational database. For visualization the embedded windows available in the Extensible Display Manager (EDM) are the EPICS clients, which retrieve interlock status information from the EPICS database. A set of interlock panels is the library, which can be used to show any chains of interlocks. If necessary, a new interlock panel can be created by using the visualization tools provided with EDM. This solution, in use for more than 3 years, has proven to be reliable and very flexible.

MOTIVATIONS

The new control systems at TRIUMF are implemented by using the Experimental Physics and Industrial Control System toolkit [1]. To control the devices in the upgraded M20 secondary beam line, Programmable Logic Controllers (PLCs) are used [2]. One of the tasks of the control system is to visualize the state of interlocks for various vacuum devices and power supplies. The present scheme for visualization of interlocks, which is used at TRIUMF was designed more than 15 years ago [3], when the old Display Manager with limited capabilities was used, and EPICS support was located in VME Input-Output Controllers (IOCs), which had only 16 Mb of memory, and 100 Megabit network connections.

At present EPICS support can be realized as an application running on a Linux server (we call it PLCIOC). This application communicates with the PLC by using the EPICS driver on one hand and with the EPICS clients on the other. Thus, PLCIOC plays the role of a transfer agent between EPICS clients and PLC: EPICS clients communicate via the network with this server, send/receive commands and set/get values from devices via PLCIOC (Fig. 1).

At present there is effectively no limitation for memory usage by PLCIOC. EPICS database can be of very large size and keep lots of information. At present we use a multi-core Dell PowerEdge 600 with 8 Gb of memory. Memory usage by M20 PLCIOC serving ~100 devices is only 1.5%. Processing information is very fast on such multi-core computers. CPU usage for M20 PLCIOC is ~12% for one core (out of 8). Transfer of data between EPICS clients and PLCIOC is also very fast for a Gigabit network connection.

Another favourable reason for new development was the addition of a new feature to the Extensible Display Manager (EDM) [4] in 2008, namely, possibility of assignment of macro value by a script or executable. Thus, it became possible to implement control device panels, in which interlocks can be represented by an embedded interlock panel, and all necessary information is retrieved from EPICS database (Fig. 2).



Figure 2: Interlock panel (on the left) is embedded in lower part of control device panel (on the right).

The goals of new implementation for interlock visualization are:

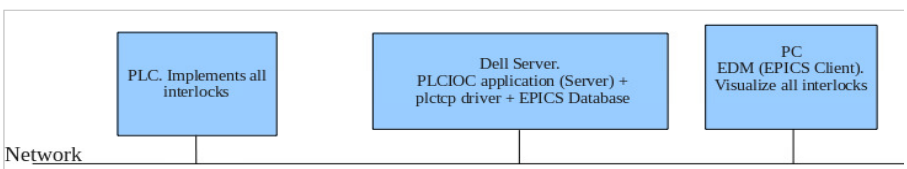


Figure1: Communication between PLC, PLCIOC, and EPICS clients.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

IMPROVING SOFTWARE SERVICES THROUGH DIAGNOSTIC AND MONITORING CAPABILITIES

P.Charrue, M.Buttner, F.Ehm, P.Jurcso, CERN, Geneva, Switzerland

Abstract

CERN's Accelerator Controls System is built upon a large set of software services which are vital for daily operations. It is important to instrument these services with sufficient diagnostic and monitoring capabilities to reduce the time to locate a problem and to enable pre-failure detection by surveillance of process internal information. The main challenges here are the diversity of programs (C/C++ and Java), real-time constraints, the distributed environment and diskless systems. This paper describes which building blocks have been developed to collect process metrics and logs, software deployment and release information and how equipment/software experts today have simple and time-saving access to them using the DIAMON console. This includes the possibility to remotely inspect the process (build-time, version, start time, counters,..) and change its log levels for more detailed information.

INTRODUCTION

Operating the CERN accelerators relies on a Controls Infrastructure in a perfect state. The CERN Beams Controls Infrastructure is deployed all around the CERN site, spanning from the Operator console in the control room to the device connected to a distant fieldbus close to the beam, via hundreds of servers, FrontEnds, networks and of course software applications.

The Beams Controls Group has studied, developed and deployed a complete diagnostic and monitoring infrastructure (DIAMON)[1] capable of detecting and repairing faults in the whole Controls infrastructure.

DIAMON

DIAMON is following the classic multi-layer system approach allowing high level of flexibility. As illustrated in Figure 1, specific DAQ modules, based on the C²MON[2] DAQ Core assure the communication between a variety of data providers and the DIAMON server. The communication can be uni- or bi-directional (support for settings and commands).

With a set of DAQ core plugins we establish communication to all required types of devices. The DIAMON server handles the messages received from the DAQ layer, computes the state of the rules (business logic) and provides the results to the client applications through C²MON Client API. The communication between components of the system layers is based on JMS messages and the scope of the monitored information is configured through a set of web applications operating on DIAMON configuration database. All configured points (further called *metrics*), alarms and rules (so-called *limits*) are handled inside the server's internal cache and the states are periodically persisted into the server's cache-persistency database. Updated metrics and limits are also periodically written into the Short-Term-Log (STL) database, which holds thirty days snapshots and provides the data on demand to the client applications. This allows our users to replay the chain of events and precisely analyse the problems retroactively. For all GUI clients and web-based client applications an instance of the Tomcat servlet container has been set-up, hosting a Spring MVC-based Java module which renders a set of web pages with the data received at runtime from the server as well as the offline data from the STL database.

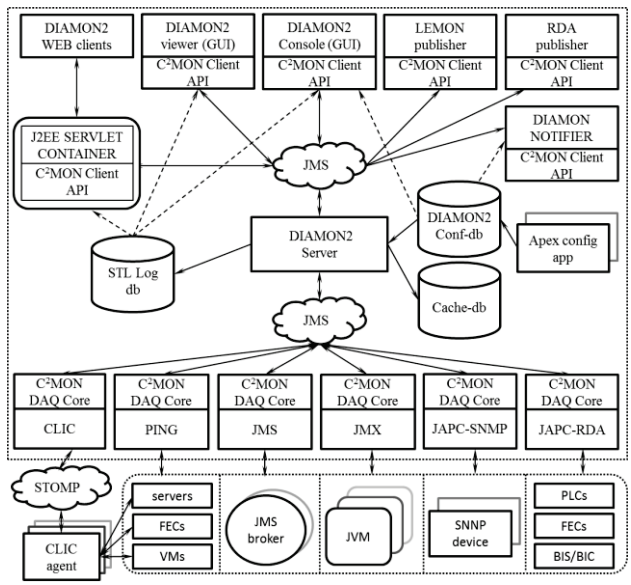


Figure 1: Overview of the DIAMON architecture.

THE DIAMON CONSOLE

The DIAMON console is the main tool that is used to retrieve and display monitoring data in a simple and homogenous way.

Specific views can be configured to target the family of devices one need to monitor, being all computers involved in the control of a specific accelerator, devices providing a specific service (such as Timing, Middleware, …), devices belonging to a specific group (such as Beam Instrumentation, RF, Power Converters, …) or the complete list of all monitored devices (~3000 at the time of writing).

Standard views are made available directly from the DIAMON console to cover the most used views, such as

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

DEVELOPING DISTRIBUTED HARD-REAL TIME SOFTWARE SYSTEMS USING FPGAS AND SOFT CORES

T. Włostowski, J. Serrano, CERN, Geneva, Switzerland
F. Vaga, University of Pavia, Italy

Abstract

Hard real-time systems guarantee by design that no deadline is ever missed. In a distributed environment such as particle accelerators, there is often the extra requirement of having diverse real-time systems synchronize to each other. Implementations on top of general purpose multitasking operating systems such as Linux generally suffer from lack of full control of the platform. On the other hand, solutions based on logic inside FPGAs can result in long development cycles. A mid-way approach is presented which allows fast software development yet guarantees full control of the timing of the execution. The solution involves using soft cores inside FPGAs, running single tasks without interrupts and without an operating system underneath. Two CERN developments are presented, both based on a unique free and open source HDL core comprising a parameterizable number of CPUs, logic to synchronize them and message queues to communicate with the local host and with remote systems. This development environment is being offered as a service to fill the gap between Linux-based solutions and full-hardware implementations.

BACKGROUND

Real-time controls in particle accelerators cover a broad range of applications that have different requirements for processing power, latency and determinism. For the purpose of this article, we divided these into the three following categories:

1. less than a microsecond of latency, such as beam injection and extraction, low level RF, interlocks and machine safety systems,
2. from dozens of microseconds to one millisecond, examples, power converter controls, orbit feedback or a number of timing and event distribution systems,
3. slow controls, such as cryogenics, vacuum, positioning or radiation monitoring.

Usually the parameter that has the highest impact on the architecture of the control system electronics is the worst case (never to be missed) processing latency. The sub-microsecond controls are the exclusive domain of FPGAs, ASICs and dedicated analog circuitry. On the contrary, the systems from point 3 are usually implemented on Programmable Logic Controllers (PLCs), or in many cases on industrial PCs running non-real time operating systems.

The widest variety of controls applications lies in between these two extremes. While the actions executed by these systems usually have strictly specified execution times, the data processing that triggers these actions can be done ahead in time, provided that the processing latency is guaranteed to never exceed a certain upper limit, as shown in Figure 1. For example, the receiver of the CERN General Machine Timing (CTR) can produce pulses with 40 nanosecond resolution, but the minimum time between reception of a timing event and the generation of a corresponding pulse is almost one millisecond.

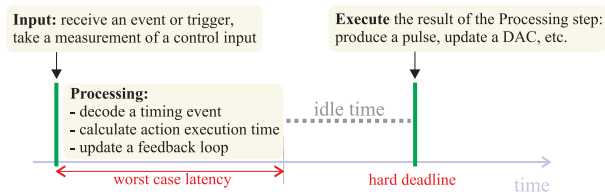


Figure 1: A typical hard-real time control system.

The traditional design approach in such systems, using an FPGA and developing custom HDL cores, brings several disadvantages:

- Only a small part of the HDL code implements the strictly timed part. In case of the CTR, these are the counters which drive the outputs and constitute around 20 % of the total HDL code. Event reception, filtering, logging and counter configuration logic has much slower execution time constraints.
- HDL development and testing take much more time compared to software development.
- Custom HDL cores need custom drivers, requiring additional development manpower.

More recent designs rely on soft processor cores, such as Xilinx MicroBlaze [1] or Altera Nios [2], which are responsible for the "slow" part, connected to custom logic implementing the "fast" part. This reduces the development time, but still leaves the following issues:

- No standard way of low-latency communication and synchronization between the devices (if they form a distributed system).
- Vendor lock-in - the cores provided by the FPGA vendors are not portable to other FPGAs,
- Issues with portability of drivers and low level host software, which are different for each FPGA manufacturer and each application.

In this article, we present the *Mock Turtle* (MT) - an HDL and software framework targeted at hard real-time distributed applications, offered as a ready to use service. *MT* provides:

- A deterministic multi-core CPU system that can be freely interfaced to user logic.
- A standardized way of communication with the host machine and (if needed) other devices in a distributed network.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

THE LANSCE FPGA EMBEDDED SIGNAL PROCESSING FRAMEWORK*

J. Hill, LANL, Los Alamos, NM 87544, USA

Abstract

During our replacement of some LANSCE LINAC instrumentation systems we have developed a common architecture for timing system synchronized embedded signal processing systems. Our design follows trends of increasing levels of electronics system integration; a single commercial-off-the-shelf (COTS) board assumes the roles of analogue-to-digital conversion and advanced signal processing while also providing the LAN attached EPICS IOC functionality. These systems are based on agile FPGA-based COTS VITA VPX boards integrating a VITA-57 FMC mezzanine site. Our signal processing is primarily developed at a high level specifying numeric algorithms in software source code to be integrated together with COTS signal processing intellectual property components for synthesis of hardware implementations. We will discuss our requirements, our decision point selecting the VPX together with the FMC industry standards, the benefits along with costs of system integrating multi-vendor COTS components, the design of some our signal processing algorithms, and the benefits along with costs of embedding the EPICS IOC within an FPGA.

SCOPE

Legacy VAX backplane integrated RICE control and instrumentation hardware systems are in service now at LANSCE for forty years. Our maintenance costs only increase and legacy schemes applying a single type of multiplexed ADC for all input signals are now obsolete. We are therefore compelled to replace them.

This paper describes electronics upgrade for a subset of RICE signals requiring real-time synchronization with a timing-system maintained data structures determining the modal operation of the facility. For example, we must flavour-stamp the incoming data streams from the sensors for correlation with the beam species being produced. The system described herein will be initially deployed within replacement loss-monitor, current-monitor, and beam-position-monitor systems at LANSCE. The embedded processor and support software described herein are utilized also within the upgraded RF control systems.

OVERARCHING REQUIREMENTS

The design is based on a reusable framework of modular off-the-shelf components. This approach offers benefits related to a shared spares, chassis-space, source-code, developer-effort, and developer-expertise. The design prefers components based on modular industry-standard long-life-span interfaces, prefers industry-standard development languages, and multiple sources for each vendor purchased modular function.

Furthermore, in a long lifespan facility, we recognize that even newly replaced systems will eventually become obsolete. With electronics components the only certainty is change. However, with a reusable framework, based on industry standard interfaces, our design can accommodate a path forward. Likewise, when one project comes to the end, if we develop our skills within a framework then hopefully some of our efforts can be preserved for use in the next project.

REQUIREMENTS

Digital Signal-Processing

Within a framework, the algorithmic requirements of individual diagnostic systems must be realizable on a shared hardware platform. In a modern system, typically an FPGA (Field Programmable Gate Array) is performing most of the signal processing tasks. Within a framework, the FPGA capacity must be sufficient to meet the needs of a diverse set of systems. Likewise, with a framework we must accommodate the signal capture requirements of many different signal types.

At LANSCE machine protection functions are currently implemented in analogue circuitry responding to inhibit the beam within a maximum overall system response of 6 μS. The signal processing framework shall be capable of responding fast enough so as to not preclude implementation of advanced machine protection functions initially supplying redundant inputs within the pre-existing machine protection system.

Real-Time Software

At LANSCE, it is necessary to beam-species flavour-stamp the incoming data streams. Due to the complexities of the modal behaviours of LANSCE, this task can't be easily accomplished in hardware, and we must therefore deploy software, running on a real-time operating system, in close proximity to the digital signal processing stream. This software must receive time deterministic interrupt response from both the timing system's event receiver module and the digital signal processing system. Certainly, this is also possible sans operating system, but an operating system is specified to improve productivity and reduce maintenance costs.

HARDWARE DESIGN

We are aware of some trends. New products offer increased density and integration; what was once possible within the confines of 6U modules is increasingly possible within the confines of 3U modules. Increasingly, a processor and Ethernet interface are cost effective for integration at the module level. We can purchase off-the-

* Supported by US Depart of Energy contract DE-AC52-06NA25396.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

MESSAGE SIGNALLED INTERRUPTS IN MIXED-MASTER CONTROL

W. W. Terpstra, GSI, Darmstadt, Germany
M. Kreider, GSI, Darmstadt, Germany

Abstract

Timing Receivers in the FAIR control system are a complex composition of multiple bus-connected components. The bus is composed of Wishbone crossbars which connect master devices to their controlled slaves. These crossbars are in turn connected in master-slave relationships forming a DAG where source nodes are masters, interior nodes are crossbars, and terminal nodes are slaves. In current designs, masters may be found at multiple levels in the composed bus. Bus masters range from embedded microcontrollers to bridges from PCIe, VME, USB, or the network.

In such a system, delivery of interrupts from controlled slaves to masters is non-trivial. The master may reside multiple levels up the hierarchy. In the case of network control, the master may be kilometres of fibre away. Our approach is to use message signalled interrupts (MSI). This is especially important as a particular slave may be controlled by different masters depending on the use-case. MSI allows the routing of interrupts via the same topology used in master-slave control. This paper explores the benefits, disadvantages, and challenges uncovered by our current implementation.

INTRODUCTION

To coordinate accelerator activities at GSI, we need a hard real-time control system. The hard real-time components of the control system are implemented using logic chips whose wiring can be reprogrammed in the field, called FPGAs. This makes it possible for our existing hardware to accommodate some of the changing requirements that arise as physicists revise their research goals. It also means that we have a great deal of flexibility in terms of how we build and connect the hard real-time components inside the FPGA.

For timing receivers, we have standardized on the pipelined Wishbone B.4 bus standard [1]. By using this open standard inside our FPGAs, we can interface components of our design more readily with components from CERN and the open hardware community at large. Our experience thus far, about six years, has shown that Wishbone strikes a good balance between simplicity and flexibility. Simple components can easily speak Wishbone, but the bus scales well enough that today our FPGA designs include nearly fifty distinct Wishbone-connected components.

In a complex system like ours, many components must communicate with each other. Most often, this communication flows from logic components, which direct the planned accelerator behaviour, to physical interfaces, which perform the external signalling to control magnets, measurement equipment, and displays. However, sometimes this communication flow must be reversed. If a magnet power supply encounters an error condition, it must be able to notify the components which control the system. The Wishbone standard does not include such a mechanism, and a traditional interrupt-based system is inadequate for our needs.

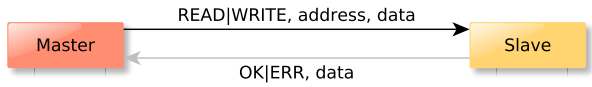


Figure 1: Wishbone connects one master to one slave. Masters initiate requests. Slaves respond with success or failure.

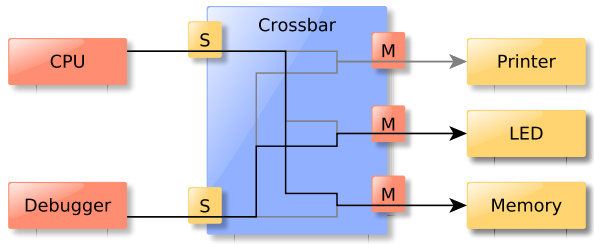


Figure 2: A crossbar providing two slave ports and three master ports. It internally connects each master to one slave.

This paper will explain how we have extended Wishbone to include a notification facility, based on message signalled interrupts (MSI). Our approach essentially mirrors the Wishbone bus to allow messages to flow bidirectionally throughout the design. While this approach carries an increased interconnect cost, it is very flexible and requires only a minor revision to the Self-Describing Bus standard [2].

WISHBONE BASICS

Wishbone is a master-slave point-to-point bus protocol. This means that it describes how to connect a single master to a single slave, as shown in Figure 1. All communication is initiated by the master, which sends bus operations (reads or writes) to the slave. The slave then executes the requested operation and reports any resulting data and the success or failure of the operation.

Bus operations include an address. Typically, the slave uses the address to determine what to do with the operation. For a slave that implements memory, the address is simply the location in its table where the value should be read or written. However, for most slaves, the address acts as an indication of the type of command. Writing to a particular address might cause the magnet to instantaneously change its field strength. Writing to a different address might cause the magnet to begin slowly ramping the field from one strength to another. Reading from one address might report the current field strength, while another address could report current or voltage. In any case, Wishbone slaves have a fixed number of addresses that a master can read or write.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

THE EVOLUTION OF THE ALICE DETECTOR CONTROL SYSTEM

P. Chochula[1], A. Augustinus[1], P. M. Bond[1], A. Kurepin[1,2], M.Lechman[1,3], O. Pinazza[1,4]
[1]CERN, Geneva, Switzerland
[2]INR RAS, Moscow, Russia
[3]Institute of Physics, Slovak Academy of Sciences, Bratislava, Slovakia
[4]INFN Bologna, Italy

Abstract

The ALICE Detector Control System has provided its service since 2007. Its operation in the past years proved that the initial design of the system fulfilled all expectations and allowed the evolution of the detectors and operational requirements to follow. In order to minimize the impact of the human factor, many procedures have been optimized and new tools have been introduced in order to allow the operator to supervise about 1 000 000 parameters from a single console. In parallel with the preparation for new runs after the LHC shutdown a prototyping for system extensions which shall be ready in 2018 has started. New detectors will require new approaches to their control and configuration. The conditions data, currently collected after each run, will be provided continuously to a farm containing 100 000 CPU cores and tens of PB of storage. In this paper the DCS design, deployed technologies, and experience gained during the 7 years of operation will be described and the initial assumptions with the current setup will be compared. The current status of the developments for the upgraded system, which will be put into operation in less than 3 years from now, will also be described.

THE ALICE DCS DESIGN

The central ALICE detector consists of 19 subdetectors, built with different detection technologies and with largely different operational requirements, all supervised by a single operator. The architecture of ALICE Detector Control System (DCS) is based on standards adopted by the LHC experiments at CERN. The commercial SCADA system WINCC OA, extended by CERN JCOP and ALICE software frameworks, is configured as a large distributed system running on about 100 servers [1]. To guarantee the autonomous operation of each subdetector, the central distributed system is segmented into subdetector systems, each allowing for stable subdetector operation in isolation from the other subdetectors.

Wherever possible, the ALICE DCS is based on commercial hardware and software standards. OPC servers handle the communication between the WINCC OA and the controls devices. Nonstandard devices, like detector frontend electronics modules, are interfaced with WINCC OA using ALICE FED standard, a complex client-server mechanism based on the DIM communication protocol [2].

All controls tasks of the ALICE experiment could be fully satisfied with the pure WINCC OA system, however its operation would require a deep knowledge of many technical details. Using the SMI++ framework [3], installed on each WINCC OA system, the operation of controlled components is modelled as a finite state machine with well-defined behaviour. Figure 1 shows the schematics of main DCS components.

The various dependencies, like the need to configure a channel before it can be turned on, are encoded in the SMI++ logic. Using the mechanisms provided in the SMI++ framework,
ALICE is represented as a hierarchical tree, with the central DCS placed on top of the pyramid and individual channels on its bottom. Commands sent from the top objects are propagated to the child nodes: the subdetectors, subsystems (high voltage, cooling, frontend electronics), devices and their channels. Each object reports its current state to its parent object. Any part of the tree can be excluded from the hierarchy and operate independently.

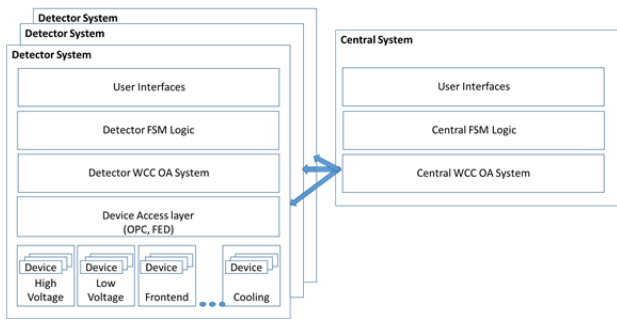


Figure 1: Architecture of the ALICE DCS.

OPERATIONAL EXPERIENCE

During eight years of ALICE operation, the initial architecture and system implementation proved their reliability and robustness. The DCS provided stable 24/7 services with small interruptions required mainly for the infrastructure modifications (cooling and ventilation upgrades, reorganization of the computer racks…). Even during the service breaks, the core systems related to safety remain operational.

With additions of new detector modules, the DCS is being continuously extended, profiting from scalability of the core architecture. In few cases new computers needed

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

STATUS OF THE CONTINUOUS MODE SCAN FOR UNDULATOR BEAMLINES AT BESSY II

A. Balzer, E. Schierle, E. Suljoti, M. Witt, HZB, Berlin, Germany
R. Follath, PSI, Villigen, Switzerland

Abstract

At the synchrotron light source BESSY II monochromator (MONO) and insertion device (ID) scans can be done synchronized in two different modes. In step mode MONO and ID move independently to intermediate target positions of an energy scan. In continuous mode (CM) MONO and ID cover the whole range of the scan nonstop in a coupled motion. Data acquisition is done continuously at the speed provided by the CM scan and is available in regular user operation [1]. Currently CM is in operation at 11 undulator beamlines at BESSY II. 3 new beamlines requesting CM are under construction. During CM the MONO EPICS IOC acts as a controller forcing the MONO optics to follow the movement of the ID. A non-linear predictive control scheme is used to implement this dynamic coupling. The controller task utilizes polynomial regression to extrapolate the ID motion. Calculation of the trajectories for MONO grating and mirror is based on bijective gap to energy lookup tables and the grating equation. In this paper the technical implementation, limitations, recently developed diagnostic methods, and future plans for improvements are presented.

INTRODUCTION

The continuous mode is part of the monochromator control program (MCCP) and has been in operation since 2005 at a dipol beamline. The first employment at an undulator beamline was 2006 [2]. At the time of writing, 11 undulator beamlines (UE56/2 PGM1, UE56/2 PGM2, UE49/1 PGM1, UE52/1 SGM1, UE52/1 PGM1, UE46/1 PGM1, UE46/1 PGM2, UE56/1 PGM1, UE112 PGM1, U49/2 PGM1, U49/2 PGM2) and 6 dipol beamlines (PM1...4, HESGM, ISISS) support CM operation by the monochromator software. 3 beamlines (U125/2 PGM1, U411/1 PGM1 and EMIL-UE48/U17 PGM [3]) are under construction requiring CM. Dipol light sources provide a continuous spectra. Energy scanning can be done by solely moving the optical elements of the monochromator which makes the implementation of CM relatively easy. On undulator beamlines the motion has to be coupled with the undulator gap and shift axes that are controlled by the insertion device control program (IDCP). The mapping between energy and gap/shift is calculated using lookup tables and piecewise polynomial or splined interpolation. The tables are held by the MCCP, which sends the target positions for gap and shift via CAN bus to the IDCP and receives the gap/shift position feedback sent from the undulator (Fig. 1). Task synchronization for combined movements has to be done by the MCCP. As a result, the monochromator acts as slave and follows the movement of the undulator (master) .

Motivation

BESSY beamlines split up into branches with up to four experimental stations per undulator light source. Depending on the application, users can greatly benefit from the fact that CM scans are about five times faster than step scans [4]. In addition, exposure times of the samples are shorter and, in some cases, the measurements are even better [5].
A robust control loop with no maintenance effort, and without the need of manual interaction, is necessary in order to ensure accurate user operation with variable scan speeds. This requirement is even more challenging for the technical setups of the beamlines currently under construction [3]. Diagnostic methods are needed to locate possible disturbances from CAN/Ethernet-jitter, inaccurate lookup tables and vibrations.

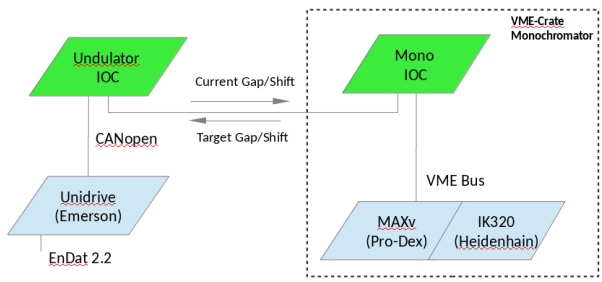


Figure 1: Interfaces used for CM from monochromator to motion controller (VME-Bus), encoder card (VME-Bus) and undulator IOC (CAN). VME based hardware for monochromator: MVME162 (Motorola), MAXv (Pro-Dex), or VME6, OMS58 (OMS), IK320 (Heidenhain).

TECHNICAL DESCRIPTION

The beamline control system is based on EPICS. Monochromator and undulator are controlled by EPICS IOCs (IDCP, MCCP). Both provide EPICS based operator interfaces. Additionally, the monochromator IOC provides connectivity over the BESSY extended monochromator control protocol (EMC) via RS232 or Ethernet. Users can control the beamline using EMC or EPICS. The communication via RS232 is slow but has proven to be reliable for CM operation when deterministic synchronization of data acquisition and monochromator energy feedback is critical. The CAN bus communication has been chosen between monochromator and undulator in order to reduce the jitter of the periodic position updates sent to the monochromator (Fig. 1).

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

MANAGING NEUTRON BEAM SCANS AT THE CANADIAN NEUTRON BEAM CENTRE

Mark Vigder, National Research Council, Canada
Martin Lee Cusick, Canadian Nuclear Laboratories, Canada
David Dean, National Research Council, Canada

Abstract

The Canadian Neutron Beam Centre (CNBC) of the Canadian Nuclear Laboratories (CNL) operate six beam lines for material research. A single beam line experiment requires scientists to acquire data as a sequence of scans that involves data acquisition at many points varying sample positions, samples, wavelength, sample environment, etc. The points at which measurements must be taken can number in the thousands with scans or their variations, having to be run multiple times. At the CNBC we have developed an approach to allow scientists to specify and manage their scans using a set of tools we have developed. Scans are specified using a set of constructors and a 'scan algebra' that allows scans to be combined using a set of scan operators. Using the operators of the algebra, long complex scan sequences can be constructed by combining shorter simpler scans. Based on the constructors and the algebra, tools are provided to scientists to build, organize and execute their scans. These tools can take the form of scripting languages, spreadsheets, or databases. This scanning technique is currently in use at CNL, and has been implemented in Python on an EPICS based control system..

INTRODUCTION

The Canadian Neutron Beam Centre (CNBC) conducts material science experiments using neutron scattering techniques. The neutron source is an experimental reactor that provides six beam lines for neutron scattering. Each beam line is specialized for a particular scanning purpose, such as triple-axis spectrometery, reflectometery, stress-scanning diffractometers, etc. In addition to the equipment controlling the spectrometer beam lines, a set of ancillary equipment is used to control the sample environment. This includes equipment such as furnaces, cryostats, stress rigs, etc.

When conducting a single experiment, scientists have a very large number of independent variables they must control. For the beam line this includes beam position and focus, energy levels on the main and diffract beams, and polarization of the main and diffract beams. Sample location and orientation must be controlled, as well as sample environment (temperature, stress, etc.) Setting each of these independent variables could involve setting multiple sub-variables, or sequencing through a number of steps to correctly set the variable.

An experiment involves setting the independent variables to the appropriate values and initiating a data acquisition phase for a specified duration. The independent variables are then reset to new values and the data acquisition repeated. A single experiment repeats this process performing a data acquisition at possibly thousands of different points within the n-dimensional space formed by the independent variables. Acquiring the data at a sequence of points in the experiment space is referred to as scanning.

Scanning is often a dynamic and interactive process. When conducting scans, scientists monitor the data as it is being acquired and will often stop, restart, modify, rerun, or resequence a scan. Creating and managing these scans and dynamically interacting with the scanning process presents a number of challenges to the scientist and to the system developer who must support the scientists work.

OBJECTIVE

The objective of this work is to develop a method that allows scientists to easily create, manage, and execute their scans. The work was undertaken during a major upgrade of the control and data acquisition system. The original spectrometer control system at NRU was developed incrementally over more than 30 years of work and ran on a VAX based system. In the course of replacing it with a distributed Linux based system, the opportunity arose to replace the outdated user interface with a modern interface, including improving the techniques for scientists to perform the scanning for experiments.

In redoing the scanning control and management, the following requirements were followed.

First, there was not to be a major paradigm shift in how the scanning process was managed by the scientists. The old system was dated, had no underlying model, and carried thirty years of baggage in its design and software. However, it also was designed by the users who knew what they needed, was very efficient and easy to learn, and carried thirty years of detailed knowledge of the scanning process. The new system was not to be a major shift in the process, but rather was to provide a better model of scanning, a refactoring of the design and implementation, and much more capability and flexibility that would be provided to users.

Second is the requirement that identical scanning control management software is to be used on all the beam lines. With minimal resources to maintain and manage the software, having different software on the different beam lines was not a viable option. Even software variants, no matter how similar, create software management problems. However because each beam line is unique this required that the system

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

PRELIMINARY DESIGN OF A REAL-TIME HARDWARE ARCHITECTURE FOR eRHIC*

R. Michnoff#, P. Cerniglia, M. Costanzo, R. Hulsart, J. Jamilkowski, W. Pekrul, Z. Sorrell, C. Theisen,
Brookhaven National Laboratory, Upton, NY 11973, USA

Abstract

The 3.8 km circumference Relativistic Heavy Ion Collider (RHIC) at BNL has been in operation since 2000. An electron-ion collider (eRHIC), which is in the design phase, plans to use one of the existing ion rings and new electron rings constructed in the existing tunnel to provide collisions of up to 21.2 GeV electrons with up to 100 GeV gold ions, 250 GeV polarized protons, as well as other species. Many new real-time systems will be required to satisfy the needs of eRHIC, including over 2000 beam position monitors, 1000 beam loss monitors, 18 current monitors, feedback systems, controls for about 10,000 power supplies, machine protection system, new beam timing systems, and more. The selected architecture must be flexible, expandable, cost-effective, reliable, and easy to maintain. Interface with existing and new accelerator timing systems is required, and compatibility with existing infrastructure and equipment must be maintained. Embedded modules based on the Xilinx Zynq gate array, with direct Ethernet connection and on-board Linux, housed in multi-slot chassis (VME, VXS, VPX, TCA, etc.) is under consideration. Preliminary design concepts for the architecture will be presented.

CONTROLS ARCHITECTURE HISTORY

Hardware architectures for accelerator controls have gone through many transformations as technology has evolved. In the 1970's, custom field-busses were designed to control remote devices in accelerators. Engineers at BNL's Alternating Gradient Synchrotron (AGS) designed a system called Datacon – a multidrop field bus capable of withstanding the electrically noisy accelerator environment [1,2]. A central computer typically controlled all of the devices. At AGS, the then state-of-the-art PDP-8 and PDP-10 were used.

As hardware evolved, real-time controls became more distributed. Intel's Multibus architecture became a standard in the 1980s for systems developed at the AGS. A typical Multibus system consisted of a custom enclosure with a backplane card cage housing processor, memory and I/O modules. A locally resident RMX deterministic real-time operating system was used. These units worked very well, but although the processor and some I/O modules were common, a major disadvantage was that each system required custom cable assemblies and connector panels to external devices. This made it very difficult to maintain spares because each system required a dedicated spare chassis.

*Work supported by Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy.
#michnoff@bnl.gov

Then in the 1990's, VME-bus became a popular platform at AGS and for the RHIC project that was under construction at that time [3]. System maintenance greatly improved with the use of VME because modules could easily be replaced when failures occurred. Swapping VME boards is very simple compared to fully encased systems like Multibus and custom designed enclosures.

More recently, manufacturers and accelerator hardware engineers are commonly incorporating direct Ethernet connectivity into systems. This has provided tremendous advantages for system designers. However, custom chassis have again become popular, which can increase the complexity of maintenance and spare inventory.

THE eRHIC APPROACH

The present plan for eRHIC real-time hardware systems is to use the Xilinx Zynq family system-on-a-chip (SoC) devices (or similar newer components that become available), which combine large programmable gate arrays with embedded ARM processors all on the same integrated circuit. A suite of modules will be developed and housed in multi-slot backplane chassis such as VME, VXS, openVPX, or TCA, with each module containing a direct Ethernet connection for communication to the higher-level control system. An operating system such as Linux will reside locally on each module.

Gbit serial links will be used for communication between modules via either front panel connections or via Gbit serial lanes on the backplane. This high-speed communication across multiple modules is essential for providing the real-time operation required for many accelerator systems. Details for the communication links are not yet fully defined, but standard protocols will be used where feasible. A draft design is provided in fig. 1.

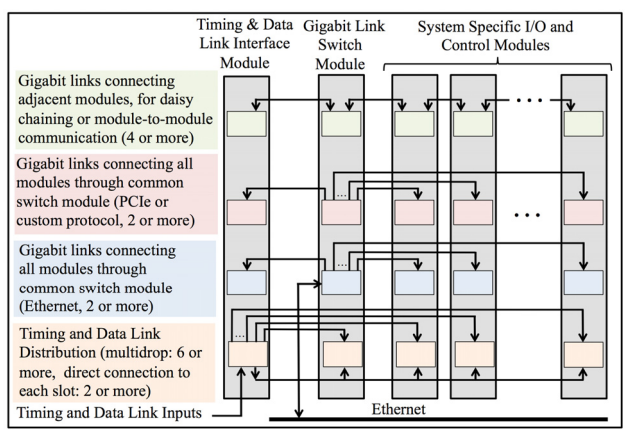


Figure 1: Draft design of inter-module communication. Communication via backplane is most desirable; some paths may be via front or rear panel connections.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

A MODULAR APPROACH TO ACQUISITION SYSTEMS FOR FUTURE CERN BEAM INSTRUMENTATION DEVELOPMENTS

A. Boccardi, M. Barros Marin, T. E. Levens, B. Szuk, W. Vigano, C. Zamantzas
CERN, Geneva, Switzerland

Abstract

This paper will present the new modular architecture adopted as a baseline by the CERN Beam Instrumentation Group for its future acquisition system developments. The main blocks of this architecture are: radiation tolerant digital front-ends, a latency deterministic multi gigabit optical link, a high pin count FMC carrier used as a VME-based back-end for data concentration and processing. Details will be given on the design criteria for each of these modules as well as examples of their use in systems currently being developed at CERN.

THE CERN BEAM INSTRUMENTATION GROUP

The Beam Instrumentation (BI) Group is responsible for designing, building and maintaining the instruments that allow observation of the particle beams and the measurement of related parameters for all CERN accelerators and transfer lines, comprising ultra-relativistic and non machines, circular and linear accelerators and even an antiproton decelerator. Each machine has very specific characteristics and requirements and therefore rarely the same instrument can be deployed on several of them. Typical beam parameters to be measured and monitored are: position, loss, intensity (total and per bunch), tune and chromaticity, luminosity, transfer lines matching, transverse and longitudinal profiles. The BI group decided to adopt a modular architecture for its acquisition system to promote the reuse of electronics and code, to simplify the maintenance and the design of this large variety of instruments.

Despite the details of each implementation can be very different the generic architecture of an instrumentation acquisition system is in the majority of the cases similar (see figure 1): the signal from the sensor is first conditioned and then digitized before being in-line processed.



Figure 1: Typical structure of an instrumentation acquisition system. The signal from the sensor is conditioned, digitized and in line processed.

For the majority of BI systems the processing is performed in a back-end unit, installed in a radiation safe area, even for the cases in which the digitalization is performed at front-end level. This allows the standardization of the back-end itself. That is one of the main block of this modular architecture, and reduce the complexity of the rad-tolerant electronics eventually required, with a consequent increase of reliability.

A problem common to all the system with a remote digitalization is the transmission of the data from the front-end to the back-end with a constant delay to avoid the need of continuous recalibrations at each reset of the system. This is particularly difficult for systems requiring large bandwidth and therefore complex serialization protocols relying on synchronization and PLLs. BI decided to adopt the multi-gigabit and latency deterministic optical link designed in the CERN PH-ESE group as the second main block of its new design toolkit, the third one being a digital front-end module implementing this link in a rad-hard version for all the systems where the digitalization stage is not only remote but also exposed to radiations.

Those three elements, on which the system specific blocks plugs, are described in the next chapters.

THE BACK-END

The board designed to be the standard back-end for the BI applications is the VFC-HD [1]. It is an FPGA-based 6U VME 64x module with a high pin count FMC (VITA 57) slot, the possibility to connect to custom VME rear transition module (RTM), 6 small form factor pluggable (SFP) slots, 4 of which dedicated to user applications and 2 to system ones, and on board DDR3.

The board was specified to be an FMC compliant carrier for its users to be able to use commercially available mezzanines. The decision to adopt the high pin count standard is linked to the need to use the latest generation of fast ADCs and DACs. Nowadays those adopt multi gigabit serial communication for their interfacing and therefore requiring the use of lines not defined in the low pin count standard of the connector.

The RTM was specified to suit the needs of systems requiring part of their connectivity to be plugged in the back of the crate.

The decision for the VME standard came from the need of being able to integrate the new board with existing systems. BI made in the past years extensive use of custom VME crates for its developments. In those crates the beam synchronous clocks and triggers are distributed by a timing receiver over the back plane and the P0

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

THE GLOBAL TRIGGER WITH ONLINE VERTEX FITTING FOR LOW ENERGY NEUTRINO RESEARCH*

Guanghua Gong, Hui Gong
Tsinghua University, Beijing, China

Abstract

Neutrino research is of great importance for particle physics, astrophysics and cosmology. A new global trigger scheme with online vertex fitting has been proposed, aiming at the ultra-low anti-neutrino energy threshold as down to 0.2MeV which is essential for the study of solar neutrino and supernova elastic scattering neutrinos on burst. With the scheme, the time of flight difference of photons fly through the liquid media from the interaction point to the surface of central detector can be corrected online, the trigger window to cover the whole spread of a specific neutrino generated photons can be significantly reduced which lessen the integrated dark noise introduced from the large amount of PMT devices hence a lower energy threshold can be achieved. The scheme is compatible, flexible and easy to implement, it can effectively help the low energy neutrino research topics.

INTRODUCTION

Neutrino research is a very popular and important topic for particle physics, astrophysics and cosmology. Several new neutrino experiments are now under construction or design such as the JUNO [1] (Jiangmen Underground Neutrino Observatory), a multi-purpose neutrino experiment, and CJPLNE [2] (China JinPing underground Laboratory Neutrino Experiment), a water-based LS neutrino experiment focusing on the solar neutrino and geo neutrino topics. The central detectors of those neutrino experiments consist large sphere or cylinder vessel to contain the giant target material which are 20kt Liquid Scintillator for JUNO and few thousand tons of water-based Liquid Scintillator for CJPLNE. PMTs are applied for photon detection due to the large sensitive area and mature technology. Around 16000 20inch PMTs will be installed for JUNO while around 15000 8inch PMTs are estimated to be deployed for CJPLNE.

A neutrino that interacts in the fiducial volume will release charge particles to further generate scintillation light. The light distribute evenly in all directions and finally be captured by the surrounding PMTs and produce electrical signals. The charge and position information of the related event could be further analysed from the PMT signals.

TRIGGER SYSTEM

The information related to one event must be collected to be distinguished by trigger system which could base on the number of PMTs that have generated signals (multiplicity trigger), or base on the total charge information (energy sum) [3]. Since the multiplicity trigger scheme requires only single hit bit information which is very easy to generate, transmit and process, it is the preferred baseline trigger method, especially for experiments that applies tens of thousands of PMTs.

For 20kt LS detector, the physical dimension will be 37meters in diameter, thus the light generated at the edge of the detector will take around 200ns time of flight before reach the PMTs installed on the opposite direction. To cover the TOF time and the slow scintillation components of the LS, the trigger system must open a 300ns time window to cover the full spread of all information from the same event.

PMT generates signal pulses due to thermionic emission which is called dark noise that can't be distinguished from the normal signal. Those dark noise will also be collected during the trigger window that could cause a coincident event. The correspond event rate is a function of trigger window width, dark noise rate and the number of PMT. Figure 1 shows the calculation result of dark noise caused coincidence rate compare to the physical event rate with parameters of 300ns trigger window, 50KHz dark noise for the new developed MCP PMT and 15000 tubes.

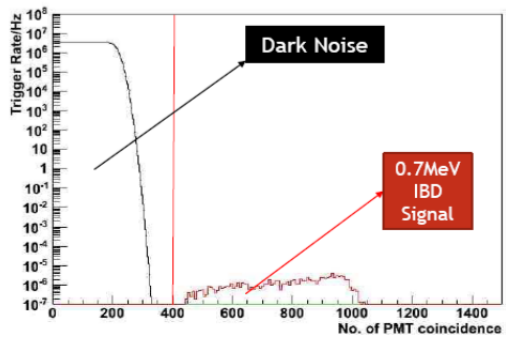


Figure 1: Event rate of PMT dark noise.

By setting the multiplicity threshold to around 400, the dark noise coincidence event can be easily separated from IBD neutrino events that have a minimum deposit energy of 0.7MeV. But for solar neutrino and supernova elastic scattering neutrino physics, many interesting phenomena are located at the energy range as low as 0.2MeV, the multiplicity distribution of such low energy event heavily overlaps with the PMT dark noise.

While approaching the low energy range, the trigger rate increases significantly. Since flash ADC is applied to

*Work supported by the National Science Foundation of China (No. 11275111), the State Key Laboratory of Particle Detection and Electronics.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

MAPPING DEVELOPMENTS AT DIAMOND

R. Walton, A.W. Ashton, M. Basham, P.C.Y. Chang, T. Cobb, S da Graca, A.J. Dent, J. Filik
M. Gerring,C. Mita, J. Mudd, C.M. Palmer, U. Pedersen, P.D.Quinn, , N. Rees
Diamond Light Source, Oxfordshire, UK

Abstract

Many synchrotron beamlines offer some form of continuous scanning for either energy scanning or sample mapping. However, this is normally done on an ad-hoc beamline by beamline basis. Diamond has recently embarked on an ambitious project to define how to implement continuous scanning as the standard way of doing virtually all mapping tasks on beamlines. The project is split into four main areas: 1) User interfaces to describe the mapping process in a scientifically relevant way, generating a scan description that can be used later; 2) The physical process of scanning and coordinating hardware motion and detector data capture across the beamline; 3) Capture of the detector data and all the associated meta-data to disk, deciding and describing the layout of the file (or files) for the main use cases; 4) Display and analysis of live data and display of processed data. In order to achieve this common approach across beamlines, the standard software used throughout the facility (Delta Tau motor controllers, EPICS, GDA and DAWN), has been built on.

INTRODUCTION

With the increase of flux obtained with most modern 3rd generation synchrotrons, collecting maps of surfaces is becoming increasingly popular. This process involves moving a sample in front of a focussed x-ray beam, then recording various interactions of this beam with the illuminated part of the sample. Interactions include, but are not limited to, fluorescence, diffraction, absorption and scattering, with each requiring specific detectors to be available on a beamline. Scans are often performed continuously as detectors are triggered and maps containing millions of individual measurements are routinely collected at specialist beamlines. The high flux available means that these scans can be performed increasingly quickly, providing significant technical challenges in all areas, from motion control and detector readout to processing and visualisation.

At Diamond Light Source there are many beamlines which can perform mapping experiments. Dealing with these challenges on a case by case basis has become unmanageable. At the time of writing we are six months into a cross beamline, cross support group project to standardise and modularise as many elements of this process as possible. As well as providing improved functionality and user experience on these beamlines a key aim is to make the solution supportable with as little effort as possible. Many aspects of the solution are applicable to techniques other than mapping (a line scan is just single row map scan) and we expect the impact of this project over the coming years to reach many of Diamond's beamlines.

To give the best chance of producing a general solution, the project targeted five beamlines with similar but unique mapping requirements. These include traditional XRF (X-ray Fluorescence) mapping [1], XRD (X-Ray Diffraction) mapping, Absorption mapping, Coherent Diffraction mapping [2] and ARPES (Angle Resolved Photoemission Spectroscopy) mapping beamlines.

At the heart of this project is a commitment to standardise the way data is handled with a new standard data flow, standardise the way continuous scans are controlled with a new EPICS control layer, and to improve the user interface and speed of collection. The remainder of this paper highlights these high level design decisions before presenting further details of the system and reporting progress.

STANDARD DATA FLOW

To simplify post processing and analysis it is important that the multidimensional and in some cases multimodal data, along with all the appropriate metadata such as experimental conditions, are stored in a standard format. Diamond uses the NeXus format [3] backed by HDF5 [4]. HDF5 performs very well, especially with the Lustre and GPFS file systems used at Diamond. Unfortunately, HDF5 cannot be read as it is written and so in order to provide live data processing and visualisation separate data paths have been maintained. To simplify this situation the HDF group was funded by Diamond, ESRF and Dectris to extend HDF5 to allow access to data as it is written. This extension is known as SWMR which stands for Single Write Multiple Read [5].

The access to live descriptive data SWMR affords greatly streamlines processing and visualisation as the same software used offline can, with some work, be used online. Further, because data can be read almost anywhere, it can be processed using tools running on a cluster which are available for post processing, allowing for sophisticated, and computationally expensive processing to be achieved in close to real time. Processed data will also be written with SWMR, providing the visualisation software live access to both processed and raw data.

MALCOLM EPICS CONTROL LAYER

The GDA data acquisition software has good support for dynamically conducting arbitrary scans, but when performance is important, such as with mapping

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

REAL-TIME DATA REDUCTION INTEGRATED INTO INSTRUMENT CONTROL SOFTWARE

P. Mutti*, F. Cecillon, C. Cocho, A. Elaazzouzi, Y. Le Goc, J. Locatelli, H. Ortiz
Institut Laue-Langevin, Grenoble, France

Abstract

The increasing complexity of the experimental activity and the growing raw dataset collected during the measurements pushed the integration of the data reduction softwares within the instrument control. On-line raw data reduction allows users to take instant decisions based on the physical quantities they are looking for. In such a way, beam time is optimise avoiding oversampling. Moreover, the datasets are more consistent and the reduction procedure, becoming now part of the sequencer workflow, is well documented and can be saved for future use. We will report on the implementation of the on-line data reduction on several instrument at the ILL as well as on the obtained performances.

INTRODUCTION

NOMAD is the instrument control software in use at the Institut Laue-Langevin. It has been designed about 10 years ago as a client/server application. The server is written in C++ to have a direct access to the C driver layer while the main client is written in Java to have a portable and reactive GUI application. A number of other client applications have been developed. Among them we can cite the Nomad Web Spy to refresh a web page that displays monitoring information, the Plot Screenshot Generator to generate offline acquisition images, etc. [1]. In the current NOMAD environment, processes and applications are running in different languages while the communication between them is based on CORBA [2]. Processes are started and stopped on demand and they can crash: - continuous development of a small team compared to the number of instruments - module-based architecture with hundred of classes The crashes are part of the problem and taken into account and may not be considered as development mistakes. It is impossible to perform the tests to ensure a 100% robustness.

The integration of reduction or more generally computation methods at the ILL supposes:

- run heterogeneous code (multiple languages e.g. Matlab, Python, etc.) owned by scientists
- run on other computers than the instrument control PC to avoid interferences
- running on different systems (Linux, Mac OS X, Windows)
- acquisition data and computation results must be exchanged in an effective way

* mutti@ill.eu

One solution is to have a monolithic centralised server that processes every computation request and sends the results asynchronously. This solution requires some extra resources to maintain the services (list of computation methods and their update, load-balancing, etc.). Another solution is to have a "microservices" approach [3]. The computation methods are distributed along the existing control and scientific computers. For that we need a fluid, flexible and easy integration. We need Erlang [4] distributed proces functionalities. In our case, those functionalities need to be integrated in C++ and Java, including multi-process, multi-environment, synchronisation and message queue as well as crash management. To be able to achieve this goal, for the development of NAPPLI we have decided to leave CORBA since it is declining technology [5].

WHAT IS NAPPLI

NAPPLI is a very lightweight application server, very easy to install and usable everywhere. NAPPLI stands for *N applications*, where N can be both interpreted as the first character of NOMAD as well as an unknown number. A NAPPLI server provides services for starting, stopping, synchronising and making distributed applications communicate. We can say that it is an application-oriented middleware. The lifecycle of remote applications can be entirely managed within the application. The server is accompanied with a client API in Java and C++ with a modern asynchronous programming model using the *future* concept [6]. The available communication patterns between the applications are request/response, publisher/subscriber (synchronised or not) and return value at the end of the execution of the application. It is possible to use the application server in a non-intrusive way. Existing applications can be called directly without using the provided API. In this case, the application itself is directly responsible of communication with the outside world. The NAPPLI services are intended to be logic and network fault tolerant. An application can terminate with an exception but the remote caller will be notified with an error, so that it will be able to take the decision to restart or not the application. The network layer also implements features to survive to failures.

Simple Example

Figure 1 describe in a simple example the way NAPPLI can be used.

The application *App1* started by NAPPLI on the computer A, requests the start of *App2* on the computer B (Fig. 1 (A)). Notice that there both the computers A and B are running NAPPLI servers. In Fig. 1 (B) once *App2* is running and

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

ON-THE-FLY SCANS FOR FAST TOMOGRAPHY AT LNLS IMAGING BEAMLINE

G. B. Z. L. Moreno[1], F. P. O'Dowd, H. H. Slepicka, R. Bongers, M. B. Cardoso
CNPEM - LNLS, Caixa Postal 6192, Campinas – 13083-970, Brazil

Abstract

As we go to brighter light sources and time resolved experiments, different approaches for executing faster scans in synchrotrons are an ever-present need. In many light sources, performing scans through a sequence of hardware triggers is the most commonly used method for synchronizing instruments and motors. Thus, in order to provide a sufficiently flexible and robust solution, the X-Ray Imaging Beamline (IMX) at the Brazilian Synchrotron Light Source [1] upgraded its scanning system to a NI PXI chassis interfacing with Galil motion controllers and EPICS environment. It currently executes point-to-point and on-the-fly scans controlled by hardware signals, fully integrated with the beamline control system under EPICS channel access protocol. Some approaches can use CS-Studio screens and automated Python scripts to create a user-friendly interface. All programming languages used in the project are easy to use and to learn, which allows high maintainability for the system delivered. The use of LNLS Hyppie platform [2, 3] also enables software modularity for better compatibility and scalability over different experimental setups and even different beamlines.

INTRODUCTION

The traditional methods for performing scans and controlling experiments at the beamline can present dead-time issues even when comparing the performance of point-to-point scans. Workarounds to the most common causes, like network latency or inefficient programming, may include the control of part of the experiment solely by hardware triggers, orchestrated by low-level implementations. At LNLS, off-the-shelf FPGA-based systems present a standard solution to controlling experiments and integrating equipment at the beamline. EPICS channel access transfers information between low-level device drivers and high-level user interfaces in the form of *EPICS Process Variables*, or EPICS PV's.

The communication between the top and bottom layers of this system architecture takes place over the network in TCP-IP protocol and thus, is non-deterministic. For fast-scan applications, one cannot disregard network latency, inherent to the communication between motion controllers, detectors, and other beamline instruments such as fast shutters and counters triggered from EPICS.

At the IMX Beamline, white beam photon density is of the order of 10^{14} ph/s/mm^2, with an energy range of 4 to 25 KeV. Hence, typical exposure times can vary from 50 to 300 ms for the most common samples. Adding this to the camera's 468 ms readout time at full resolution, frame period could reach more than 500 ms. Due to network overhead while interpreting EPICS PVs, frame period can reach almost a second in a conventional scan. The hardware synchronization is a way to eliminate this latency given the fact that, after its start, the scan runs independently under low-level applications, not relying on communication over the network.

[1]gabriel.moreno@lnls

SYSTEM DESCRIPTION AND DATA ACQUISITION ARCHITECTURE

The LNLS X-Ray Imaging Beamline (IMX) uses a National Instruments PXI-6602 timing board attached to an NI PXI-1045 chassis to read digital counters and trigger devices during scans. A Huber 409 rotation stage performs the sample rotation during tomography. The scan points are programmed directly on Galil DMC-4183 motion controllers as a function of parameters passed from EPICS PV's to reduce latency between scan points and enable quick tomography. The controller sends feedback through the I/O interface on trippoint arrival and transmits the present motion status, enabling hardware synchronization between motors and scan devices. A LabVIEW VI running on a dedicated Windows Machine controls a PCO.2000 camera, which is currently the main detector at IMX Beamline. To prevent high x-ray dose on the sample, a Uniblitz XRS25 shutter blocks the x-ray outside acquisition interval. The user interface screens run from a workstation on CS-Studio environment under Linux Red Hat operational system. CS-Studio inputs integrate to EPICS via Py4Syn scripts to run scan routines and set up experiments.

Simplifying the experiment control system to a three-layered architecture, the EPICS channel access protocol works as an intermediate layer (or "*Service Layer*"). It transfers information between low-level device drivers and the high-level user interface. The graphical user interfaces and experiment command sequences stay on the top layer, or so-called *"Application Layer"*, assigning parameters in the form of EPICS PV's to EPICS IOC's (*Input-Output Controllers*). These IOC's on the "*Service Layer*" translate the parameters received from above into commands and signals to the device drivers on the layer below, the "*Driver Layer"*. To connect the PXI's FPGA environment to EPICS Records, an in-house data exchange protocol called Hyppie [2, 3] creates a bridge between EPICS and the low-level hardware control. The FPGA side runs on LabVIEW Real-Time, as the EPICS side runs on Red Hat Linux. Arriving commands from the application layer land at the Linux side, passing to the RT side over shared memory on the PXI chassis. The hard-

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

BEAM PROPERTY MANAGEMENT AT KEK ELECTRON / POSITRON 7-GeV INJECTOR LINAC

K. Furukawa*, N. Iida, T. Kamitani, S. Kazama, T. Miura, F. Miyahara, Y. Ohnishi, M. Satoh, T. Suwada, K. Yokoyama, KEK, 1-1 Oho, Tsukuba, Ibaraki, 305-0801, Japan

Abstract

The electron / positron injector linac at KEK has injected a variety of beams into the electron accelerator complex of an asymmetric collider and light sources for particle physics and photon science experiments for more than 30 years. The beam property of electrons and positrons varies in energy from 2.5 GeV to 7 GeV and in bunch charge from 0.2 nC to 10 nC, and their beam emittance and stability requirements are challenging dependent on the injected storage rings. They have to be switched by pulse-to-pulse modulation at 50 Hz. The emittance control is especially crucial to achieve the goal at SuperKEKB and is under development. The beam energy management becomes more important as it affects all of the beam properties. Beam acceleration provided by 60 high-power microwave stations should be properly arranged considering redundancy and stability. Thus, the equipment controls are also restructured in order to enable the precise control of the beam properties, based on the synchronized event control system and EPICS control system. The strategy and status of the upgrade is discussed from the practical aspects of device controls, online simulation and operation.

INTRODUCTION

At High Energy Accelerator Research Organization (KEK), SuperKEKB, the electron-positron asymmetric collider, is under construction [1]. This project is expected to be able to elucidate the flavor physics of elementary particles with 40-fold improved luminosity compared with the initial KEKB, by doubling the stored beam current, and also by the nano-beam scheme to shrink the beam size down to a twentieth at the interaction point. SuperKEKB is composed of the electron-positron injector, high-energy electron ring (HER), low-energy positron ring (LER), and a positron damping ring (DR).

The Electron-positron injector will perform the first injection to SuperKEKB in Japanese fiscal year 2015. It will fill the dual rings of SuperKEKB as well as two light source rings in top-up mode in 2017 with significantly different beam properties while switching beams at 50 Hz as shown in Fig. 1. The beam energies are 2.5 GeV for Photon Factory (PF), 6.5 GeV for PF Advanced Ring (PF-AR), 4.0 GeV positron for LER and 7.0 GeV electron for HER, respectively [2].

Especially, it is a major challenge to inject beams to SuperKEKB with a small emittance of 20 mm·mrad and an energy-spread of 0.1% under the large beam current of 5 nC per bunch. In order to achieve such a high quality beam, a new RF gun for high-current and low-emittance electron beam and a flux concentrator for high-current positron capture are introduced. It is crucial for the injection operation with higher stability and accuracy to achieve lateral equipment alignment less than 0.3 mm over 600 m linac, and less than 0.1 mm in a short segment. The designed beam should be delivered with iterative corrections understanding the static and dynamic properties of the accelerator equipment and evaluating the beam properties in every conceivable way.

These beam operation management processes need to be improved for SuperKEKB on the basis of the techniques achieved in KEKB.

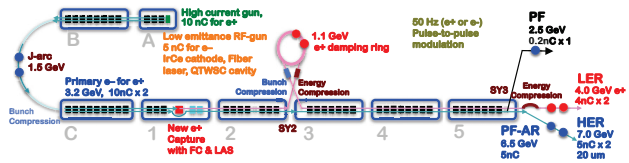


Figure 1: Layout of injector linac and its beam delivery to multiple experimental facilities.

ACCELERATOR EQUIPMENT

The injector linac is operated with the EPICS control framework at the lower and middle layer [3], the event-based control system (MRF) at the lowest layer, and the script languages including SADscript for online accelerator design. This combination was quite successful at the both KEKB injector and collider rings, and is maintained for SuperKEKB as well, with many improvements such as embedded EPICS systems [4,5]. This environment is supported by the carefully-managed control databases.

Example of static database of accelerator equipment for beam-property management is shown in Fig. 2.

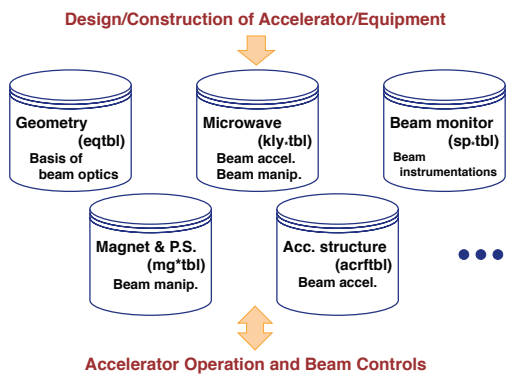


Figure 2: Construction of static database and its contribution to linac beam operation.

* kazuro.furukawa@kek.jp

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

COMPONENT DATABASE FOR THE APS UPGRADE*

S. Veseli, N.D. Arnold, D.P. Jarosz, J. Carwardine, G. Decker, N. Schwarz, Argonne National Laboratory, Argonne, IL 60439, USA

Abstract

The Advanced Photon Source Upgrade (APS-U) project will replace the existing APS storage ring with a multi-bend achromat (MBA) lattice to provide extreme transverse coherence and extreme brightness x-rays to its users. As the time to replace the existing storage ring accelerator is of critical concern, an aggressive one-year removal/installation/testing period is being planned. To aid in the management of the thousands of components to be installed in such a short time, the Component Database (CDB) application is being developed with the purpose to identify, document, track, locate, and organize components in a central database. Three major domains are being addressed: Component definitions (which together make up an exhaustive "Component Catalog"), Designs (groupings of components to create subsystems), and Component Instances (“Inventory”). Relationships between the major domains offer additional "system knowledge" to be captured that will be leveraged with future tools and applications. It is imperative to provide sub-system engineers with a functional application early in the machine design cycle. Topics discussed in this paper include the initial design and deployment of CDB, as well as future development plans.

OVERVIEW

The Component Database (CDB) application is a tool for organizing and tracking components and designs used for the APS storage ring upgrade. It helps capture component documentation, provides a repository for inspection and measurement data (e.g., travellers), and supports logging of component history through the component’s life cycle.

CDB also serves as a user portal for finding all known information about a particular component or a design. To that end, it provides links and interfaces to external systems commonly used at APS, such as various drawing and document management systems, procurement applications, etc.

Although CDB has been designed and developed from the ground up in order to satisfy APS-U requirements, in many respects it draws ideas from IRMIS2 [1], which is still in use by Controls Group at APS.

SOFTWARE COMPONENTS

CDB is built around relational database, web portal and REST web service [2] technologies. The architecture, shown in Fig. 1, provides users with a number of options for accessing the system. At the same time, it also offers developers a fair amount of flexibility for integration with external applications. The most important CDB system components are described below in more details.

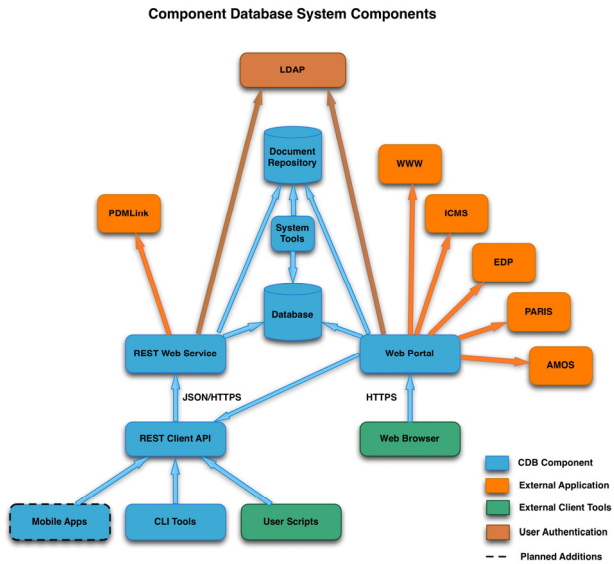


Figure 1: CDB system architecture. Dashed lines indicate future components.

Relational Database

The database contains all system data, other than uploaded documents. Users cannot access the database directly. Except for a small number of administrative tools, other software components access the database through object-relational mapper (ORM) libraries and APIs, which provide an abstraction layer from the rest of the system. CDB currently uses a MySQL database [3].

Document Repository

The Document Repository is a storage area designated for use and managed by the CDB software. It stores various documents, images, and log attachments uploaded by CDB users.

Web Portal

The User Web Portal is a Java EE application running in a GlassFish application server [4], and built using modern technologies like Java Persistence API (JPA) and Java Server Faces (JSF) [5]. In particular, CDB uses the PrimeFaces component suite [6].

Web Service

A REST Web Service and its Client APIs provide programmatic interfaces for accessing the system that are based on JSON data-interchange format [7] over the

* Argonne National Laboratory's work was supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under contract DE-AC02-06CH11357.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

REPLACING THE ENGINE IN YOUR CAR WHILE YOU ARE STILL DRIVING IT*

E.Bjorklund#, LANL, Los Alamos, NM 87545, USA

Abstract

Replacing your accelerator's timing system with a completely different architecture is not something that happens very often. Perhaps even rarer is the requirement that the replacement not interfere with the accelerator's normal operational cycle.

In 2011, The Los Alamos Neutron Science Center (LANSCE) began the purchasing and installation phase of a nine-year rolling upgrade project which will eventually result in the complete replacement of the low-level RF system, the timing system, the industrial I/O system, the beam-synchronized data acquisition system, the fast-protect reporting system, and much of the diagnostic equipment [1]. These projects are mostly independent of each other, with their own installation schedules, priorities, and time-lines. All of them, however, must interface with the timing system.

INTRODUCTION

LANSCE had its beginning in 1972 as an 800 MeV "meson factory" [2]. Since then it has expanded its missions to include such diverse projects as pion treatment for inoperable cancers, spallation neutrons, ultra-cold neutrons, medical isotope production, and proton radiography.

In preparation for the new MaRIE project [3], we are undertaking an ambitious overhaul of a large part of the facility while still trying to maintain a viable user program.

This paper will focus mostly on the timing system replacement project, its conversion from a home-built, centralized, discrete signal distribution system, to a commercial event-driven system from Micro Research Finland [4]. We will explore some of the challenges faced by having to interface with both the old and new equipment until the upgrade is completed.

PROJECT STRATEGY AND SCHEDULE

The installation/operations schedule can be compared to driving through mountainous terrain on a road with many peaks and valleys. When you start down a valley, you shut down your engine, replace as much of it as you can, then try to get it running again before you have to start up the next peak. At the bottom of each valley there is a relatively flat stretch of road representing the "startup period" – during which you mostly coast while you discover how your changes affected the machine's operation (for good or for ill).

The current installation and operation schedule is shown below in Figure 1. The green blocks represent the operating periods, the red blocks represent the installation and maintenance periods, and the yellow blocks represent the startup periods. The durations of the operation, maintenance, and startup periods vary as the project progresses. The first three years of the schedule call for longer operational periods (seven to nine months), shorter upgrade periods (four months), and shorter startup periods (one month). The middle three years – during which the most complex upgrades take place – have longer maintenance periods (four to five months), longer startup periods (three months), and shorter operational periods (three to four months). During the last three years, things theoretically get easier and we go back to longer operations, shorter maintenance, and shorter startup periods.

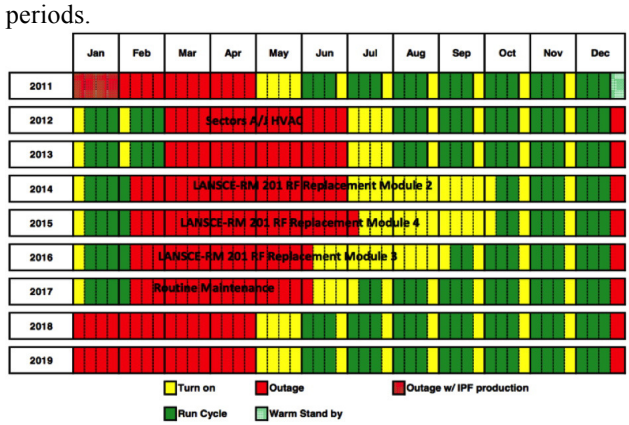


Figure 1: Installation And Operation Schedule

Budget Schedule

Our controls group adopted a budgeting strategy of purchasing all the equipment for each project at once. With some exceptions, for each fiscal year a different project had its own "year of profligate spending". The first year it was the network installation. The second year it was the industrial I/O system. Next was timing, etc. There were a number of reasons for adopting this strategy.

One technical reason was uniformity. All of the equipment for a sub-system would be from the same vendor, with the same firmware level, and therefore have a uniform interface to the controls software.

A scheduling reason was flexibility. Once you have had "your year" your system doesn't have to worry about having enough equipment to meet other projects' sometimes unpredictable schedules. This, of course, implies that those sub-systems most depended on by other sub-systems (for example, network and timing) will need to be financed earlier.

One financial reason was an uncertain funding profile (as described in [1]). Purchasing everything at once

*Work supported by US DOE under contract DE-AC52-06NA25396
bjo@lanl.gov

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

THE MeerKAT GRAPHICAL USER INTERFACE TECHNOLOGY STACK

M. Alberts*, SKA SA, Cape Town, South Africa
F. Joubert#, SKA SA, Cape Town, South Africa

Abstract

The South African MeerKAT radio telescope, currently being built some 90 km outside the small Northern Cape town of Carnarvon, is a precursor to the Square Kilometre Array (SKA) telescope and will be integrated into the mid-frequency component of SKA Phase 1. Providing the graphical user interface (GUI) for MeerKAT required a reassessment of currently employed technologies with a strong focus on leveraging modern user interface technologies and design techniques. An extensive investigation was performed to evaluate and assess potential GUI technologies and frameworks. The result of this investigative study identified a responsive web application for the frontend and asynchronous web server for the backend. In particular the AngularJS framework used in combination with Material Design principles, Websockets and other popular javascript layout and imaging libraries, such as D3.js, proved an ideal fit for the requirements of the MeerKAT GUI frontend. This paper will provide a summary of the user interface technology investigation and further expound on the whole technology stack adopted to provide a modern user interface with real time capabilities.

INTRODUCTION

MeerKAT is a mid-frequency "pathfinder" radio telescope and precursor to building the world's largest and most sensitive radio telescope, the Square Kilometre Array (SKA). MeerKAT builds upon its own precursor, KAT-7 (Karoo Array Telescope), a seven-dish array currently being used as an engineering and science prototype.

During the preliminary stages of the MeerKAT control and monitoring design it became evident that the current KAT-7 user interface and its underlying technologies have various shortcomings and will not scale well. This prompted an investigation into modern user interface technologies, especially web frameworks with all its accompanied benefits.

This paper starts with a brief description of the design methodology adopted. Next, this paper provides a summary of the user interface technology investigation. As main focus, this paper presents the new MeerKAT GUI architecture with detail on the architecture and technology stack.

METHOD

High level requirements for the MeerKAT user interface were defined by System Engineering and further refined through bi-monthly discussions with the relevant stakeholders, especially the telescope operators and commissioners. During these meetings all requirements were clarified, additional operator requirements were defined, and mock-up displays were drawn to ensure the resulting interface will fulfill all the needs of the end users.

Following an iterative development approach, combined with monthly demonstrations of prototype displays to relevant stakeholders, we were able to obtain valuable feedback that we could include in the development cycle.

TECHNOLOGY INVESTIGATION

After conducting a research exercise into GUI technologies for Responsive Web Design (RWD), various frameworks and libraries were identified. Categorising these GUI technologies according to their main function showed that few were all-in-one solutions for a user interface development platform. Most of them focus on a specific area of user interface design and should be used in conjunction with other libraries to construct a complete frontend development platform solution.

To overcome analysis paralysis we've limited our evaluation to only a select few technology solutions based on industry popularity and peer recommendations, namely AngularJS [1], EmberJS [2] and CS-Studio BOY [3].

With the help of the Control and Monitoring (CAM) team we identified criteria to evaluate the chosen technologies against while building prototypes using each technology as a frontend design platform. The prototypes were based on a typical antenna control and monitoring use case. The evaluation criteria included open source licensing, MVW (Model View Whatever) framework, good documentation, large example base, large active developer community, large scientific user community, pre-built standard widgets, in-house knowledge, user-defined widgets, rapid prototyping toolset, template system, support for testing frameworks, deployment/server side simplicity, easy learning curve, security support (authentication/authorisation), Python support, flexible data-bindings and flexible layouts.

Final scoring of the three prototyped user interface technologies made it clear that web technologies are best suited to our needs, with AngularJS coming out as the favourite.

MEERKAT GUI ARCHITECTURE

Architecture Design Overview

On a high level, the MeerKAT user interface implements a client-server architecture. A frontend component provides the client-side functionality and a

* talberts@ska.ac.za
fjoubert@ska.ac.za

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

EFFORTLESS CREATION OF CONTROL & DATA ACQUISITION GRAPHICAL USER INTERFACES WITH TAURUS

C. Pascual-Izarra#, G. Cuní, C. Falcón-Torres, D. Fernández-Carreiras, Z. Reszela, M. Rosanes,
ALBA-CELLS Synchrotron, Cerdanyola del Vallès, Spain
T. Coutinho, ESRF, Grenoble, France

Abstract

Creating and supporting Graphical User Interfaces (GUIs) for experiment control and data acquisition has traditionally been a major drain of time and resources for laboratories. GUIs often need to be adapted to new equipment or methods, but typical users lack the technical skills to perform the required modifications, let alone to create new GUIs. Here we present the Taurus [1] framework which allows a non-programmer to create a fully-featured GUI (with forms, plots, synoptics, etc.) from scratch in a few minutes using a "wizard" as well as to customize and expand it by drag-and-dropping elements around at execution time. Moreover, Taurus also gives full control to more advanced users to access, create and customize a GUI programmatically using Python [2]. Taurus is a free, open source, multi-platform pure Python module (it uses PyQt [3] for the GUI). Its support and development are driven by an active and welcoming community participated by several major laboratories and companies which use it for their developments. While Taurus was originally designed within the Sardana [4] suite for the Tango [5] control system, now it can also support other control systems (even simultaneously) via plug-ins.

INTRODUCTION

Taurus is a framework for creating user interfaces (both GUIs and command-line based) to interact with scientific and industrial control systems as well as with other related data sources.

In this work we first give a brief overview of Taurus and then we focus on one of its key assets: the possibility of deploying powerful, customizable and flexible control and data acquisition GUIs within minutes without requiring to write a single line of code. Finally, we discuss the current development efforts and the future plans.

OVERVIEW

Background

Taurus was originally conceived (under early internal names such as Tauico, Tauiwi, and Tau) as the ALBA [6] synchrotron's in-house solution for connecting client side applications to Tango device servers [7]. It provided the user interface code for the Sardana suite [8], which is used in ALBA for control and data acquisition of both the accelerator and all the beamlines.

After its first public release in 2011, Taurus has been adopted by several large laboratories and companies, and it has become very popular among many newcomers to the Tango Collaboration[1].

Since 2012 much effort has been put into bringing control system agnosticity into Taurus: proof-of-concept plugins for supporting EPICS [9] and SPEC [10] were developed and the Taurus core has been refactored in order to isolate the Tango dependencies into an optional plugin for the next major release (Taurus 4)[2].

Community

The Taurus Community largely intersects with the Sardana Community (from which it spun off to reflect the fact that Taurus can be used independently of Sardana) and shares its organizational characteristics and open development model [11, 12]: public code review process, proposal-driven decision taking [13], periodical meetings (~yearly), free licensing (LGPLv3+ [14]), etc.

FAST GUI CREATION

Model-View-Controller Approach

Taurus uses the Model-View-Controller (MVC) pattern [15] to build interfaces[3].

The *taurus.core* module uses plugins (known as *schemes*) to provide *TaurusModel* objects that abstract the interactions with specific sources of data and/or control objects. Some schemes are already implemented for accessing control system libraries (the “tango”, “epics” and “spec” schemes) as well as for data-processing via a Python interpreter (the “*evaluation*” scheme).

Every TaurusModel object can be of type *Authority*, *Device* or *Attribute*, and has a unique name in the form of a Unified Resource Identifier (URI). See Table 1 for some examples of model names. Each scheme implements a

#cpascual@cells.es

[1] While early adopters of Taurus were mostly synchrotrons (*Desy, MAX-IV, Solaris, ESRF,...)* other scientific institutions such as the *ELI-ALPS* and *LULI-APOLLON* large laser installations or the *ONERA* wind tunnel are currently using it. Companies such as *Cosylab*, *Nexeya*, *Tata Consultancy Services* and *Observatory Sciences are* providing services based on Taurus to a growing number of institutions, including the world largest radio telescope (SKA).

[2] Unless explicitly stated otherwise, in the rest of the article the situation described corresponds to the current Taurus4 state.

[3] Many Taurus components combine the View and Controller roles. Those cases could be referred to as “Model-View” instead of MVC.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

NATIONAL IGNITION FACILITY (NIF) EXPERIMENT INTERFACE CONSOLIDATION AND SIMPLIFICATION TO SUPPORT OPERATIONAL USER FACILITY

A. Casey, E. Bond, B. Conrad, M. Hutton, P. Reisdorf , S. Reisdorf.
Lawrence Livermore National Laboratory, Livermore, CA 94550 USA

Abstract

The National Ignition Facility (NIF) at the Lawrence Livermore National Laboratory is a 192-beam 1.8 MJ ultraviolet laser system designed to support high-energy-density science. NIF can create extreme states of matter, including temperatures of 100 million degrees and pressures that exceed 100 billion times Earth's atmosphere. At these temperatures and pressures, scientists explore the physics of planetary interiors, supernovae and thermonuclear burn. In the past year, NIF has transitioned to an operational facility and significant focus has been placed on how the users interact with the tools necessary to conduct an experiment at NIF. The current toolset was developed with a view to commissioning the NIF and thus allows flexibility that most users do not require. The goals of this effort include enhancing NIF's external website presence, easier proposal entry for NIF experiments, reducing both the amount and frequency of data the users have to enter, and simplifying user interactions with the tools while reducing the reliance on custom software. This paper will discuss the strategies adopted to meet the goals, highlight some of the user tool improvements that have been implemented and planned future directions for the toolset.

INTRODUCTION

In the past year, the National Ignition Facility (NIF) [1] has transitioned to an operational User Facility and the tools accessed by users needed to be updated accordingly. The current tools were designed to support facility commissioning and have the following characteristics:

- Designed for commissioning a specific function and are integrated with other tools after the fact
- Provide maximum flexibility and options over simplicity
- Tool access is assumed to be from on-site
- Users assumed to have expert level knowledge and experience with the tools

To get an idea of what this looks like to the user, Figure 1 lists the tools that are be used to define an experimental setup on NIF. Even for the expert user, this is a significant number of tools and hence options that the user has to understand. This ultimately detracts from the real task of defining and executing experiments.



Figure 1: Users Need to Interact with Several Experiment Setup Tools at NIF

The goal of the team was to make interacting with NIF analogous to buying an airplane ticket, Figure 2. The user of the airline website needs to only know some high level details of what they want to accomplish such as departure and return dates, number of people in the party etc. They do not need to know about plane weight, fuel load balances, airport operating hours etc. that are needed by the airlines in order to operate safely and get the customer to their chosen destination. In the same way, to conduct an experiment, the user should only need to provide the information relevant to the task that they are trying to complete.

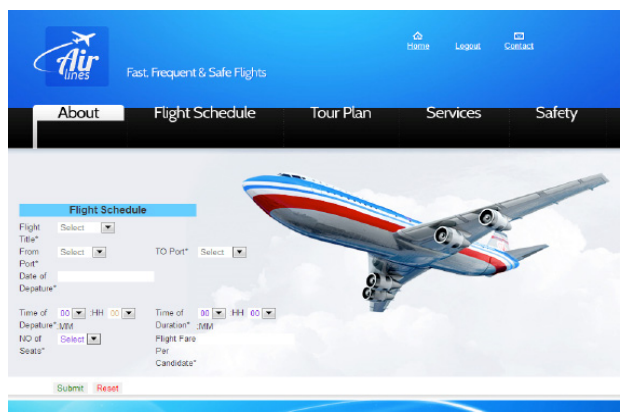


Figure 2: A web site that supports simple user interaction.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

CONTROL SYSTEMS FOR SPALLATION TARGET IN CHINA INITIATIVE ACCELERATOR DRIVEN SYSTEM*

Zhiyong He*, Qiang Zhao,Wenjuan Cui,Yuxi Luo, Ting Xie, Xueying Zhang, Lei Yang, Hushan Xu
Institute of Modern Physics, Chinese Academy of Sciences, Lanzhou, China

Abstract

In this paper, we report the design of the control system for the spallation target in China initiative accelerator driven sub-critical (ADS) system, where a heavy-metal target located vertically at the centre of a sub-critical reactor core is bombarded vertically by the high-energy protons from an accelerator. The main functions of the control system for the target are to monitor and control thermal hydraulic, neutron flux, and accelerator-target interface. The first function is to control the components in the primary and secondary loops, such as pumps, heat exchangers, valves, sensors, etc. For the commissioning measurements of the accelerator, the second function is to monitor the neutrons from the spallation target. The three-layer architecture has been used in the control system. In the middle network layer, in order to increase the network reliability, the redundant Ethernet based on Ethernet ring protection protocol has been considered. In the bottom equipment layer, the equipment controls for the above-mentioned functions have been designed. Finally, because the main objective of the target is to integrate the accelerator and the reactor into one system, the integration of accelerator's control system and the reactor's instrumentation and controls into the target's control system has been mentioned.

INTRODUCTION

Driven by the national demand for safe disposal of nuclear waste as well as the potentials for advanced power generation, the Chinese Academy of Sciences initiated an accelerator driven sub-critical (ADS) program in 2011 under the frame of "Strategic Priority Research Program". The ultimate goal of the China ADS program is to build an industrial demo facility for ADS technology. In the China ADS system, a heavy metal spallation target located at the centre of a sub-critical core is bombarded by the high-energy protons from an accelerator. An overall control system is required to exactly couple the high-energy beam from the accelerator to the spallation target in the reactor core, by controlling and coordinating three facilities, i.e. an accelerator, a spallation target and a reactor.

In this paper, we report the design of control system for the spallation target. When designing this control system, the following two objectives have to meet. The first objective is to control the demo facility of spallation target during the commissioning of the target without proton beam. The commissioning without beam may be performed either within the reactor core or outside the core. The second objective is to control the spallation target as a part of the China ADS system, when three facilities, the accelerator, the spallation target and the reactor, have been integrated into one system. Therefore, the same architecture as the overall control system should be used in the control system for the spallation target.

* corresponding author, Email: zyhe@impcas.ac.cn
This work is supported by Strategic Priority Research Program of Chinese Academy of Sciences (XDA03010000 and XDA03030000).

ARCHITECTURE OF CONTROL SYSTEM

The spallation target used in the China ADS system includes several sub-systems. According to the locations of the devices, the subsystems can be classified as the target core subsystem, the primary cooling loop, the secondary cooling loop, the target window and its cooling loop, fill and drain subsystem, cover gas subsystem, and other auxiliary subsystems. The control system for the spallation target controls and coordinates these subsystems during startup, ascent to power, power operation, and shutdown conditions of China ADS system. The control system provides automatic as well as manual controls, monitoring and diagnostics capabilities for these subsystems. Furthermore, since the spallation target is a critical part of the ADS system, its control system should be able to be integrated into the overall control system for the China ADS system.

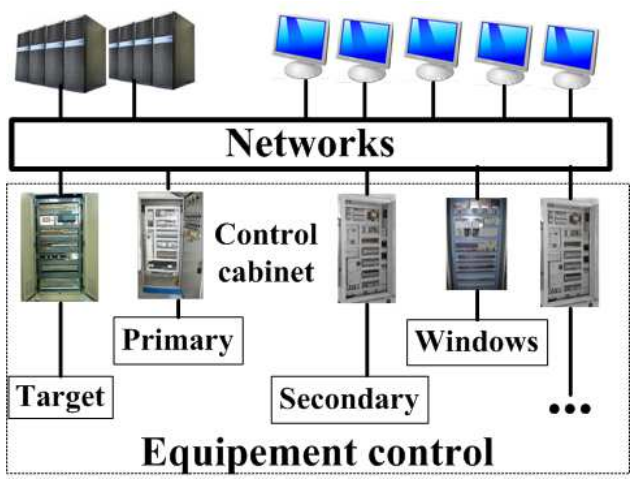


Figure 1: The three-layer architecture of the control system for the spallation target. The legends target, primary, secondary, and window represent the devices for these subsystems.

As shown in Fig. 1, the three-layer architecture is used in the control system, i.e. the top operation layer for the

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

STANDARDS-BASED OPEN-SOURCE PLC DIAGNOSTICS MONITORING

B. Copy, M. Zimny, H. Milcent
CERN, Geneva, Switzerland

Abstract

CERN employs a large number of Programmable Logic Controllers (PLCs) to implement industrial processes. These PLCs provide critical functions and must be placed under permanent monitoring. However, owing to their proprietary architecture, it is difficult to both monitor the status of these automates using vendor-provided software packages and integrate the resulting data with the CERN accelerator infrastructure, which itself relies on CERN-specific protocols and configuration facilities.

This paper exposes the architecture of a stand-alone "PLC diagnostics monitoring" Linux daemon, which provides live diagnostics information through standard means and protocols, namely file logging, CERN protocols, and Java Management Extensions. Such information is currently consumed by the EN-ICE MOON supervision software [1] used by the EN-ICE Standby Service to monitor the status of critical industrial applications used in the LHC and the CERN DIAMoN monitoring console [2] used by the LHC operators. Both applications are used daily to monitor and diagnose critical PLC hardware running all over CERN.

PLC DEVICES AND THE RELEVANCE OF INDUSTRIAL STANDARDS

PLCs are off-the-shelf industrial components designed to operate in a sheltered network environment. Moderate, predictable network traffic to and from the equipment and a small number of concurrent network-based accesses (in order to perform control with near real-time precision) are both essential parameters for ensuring reliable behavior on the part of the controller.

Many such controllers are in operation inside the LHC complex. While PLC device implementations are typically subject to IEC standards, ensuring that code is somewhat portable between vendors, there are no such standards for PLC diagnostics. In order to be able to query the internal status of a PLC (such as its last known cycle time, its operational mode or the current usage of network resources), it is necessary to resort to mechanisms specific to each PLC vendor. However, owing to the large number of PLCs deployed at CERN, we cannot simply monitor them and configure their location one-by-one; we must have a common, vendor-independent way of identifying the PLCs under monitoring.

Finally, the data resulting from diagnostic calls made to PLCs is also non-standard. It must be interpreted, transformed to a standard data format, and then stored if it is to be integrated with an established infrastructure monitoring system, such as the EN-ICE MOON platform or the CERN DIAMoN console.

To summarize, obtaining diagnostics information from a PLC is subject to three requirements:

- **Requirement 1:** To minimize as much as possible the impact of status monitoring on the PLC itself and impose absolutely no changes to the program the PLC is running (for instance, we cannot require that the PLC asset owner implements a diagnostics routine for us; it must remain completely transparent).
- **Requirement 2:** To be able to configure the inventory of PLCs to be monitored in a vendor-independent manner.
- **Requirement 3:** To bridge between the vendor-specific diagnostic data formats and a centralized device status history visualization facility.

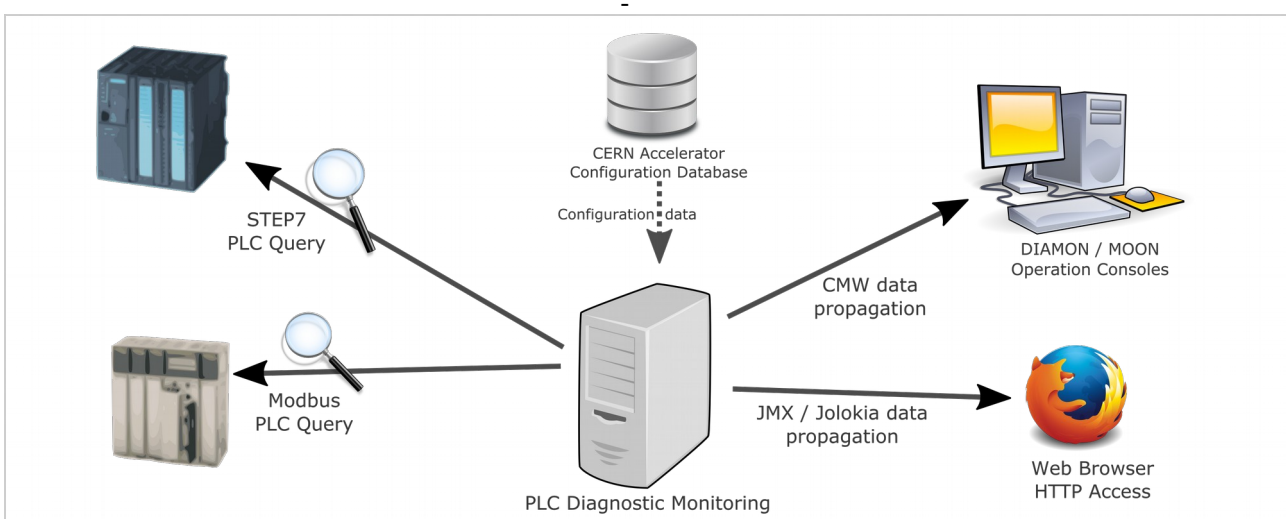


Figure 1: PLC Diagnostic Monitoring.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

OVERVIEW OF THE MONITORING DATA ARCHIVE USED ON MeerKAT

M. Slabber*, SKA SA, Cape Town, South Africa

Abstract

MeerKAT [1], the 64-receptor radio telescope being built in the Karoo, South Africa, by Square Kilometre Array South Africa (SKA SA), comprises a large number of components. All components are interfaced to the Control and Monitoring (CAM) system via the Karoo Array Telescope Communication Protocol (KATCP). KATCP is used extensively for internal communications between CAM components and other subsystems [2]. A KATCP interface exposes requests and sensors [3]. Sampling strategies are set on sensors, ranging from several updates per second to infrequent updates. The sensor samples are of multiple types, from small integers to text fields. As the various components react to user input and sensor samples, the samples with timestamps need to be permanently stored and made available for scientists, engineers and operators to query and analyse. This paper present how the storage infrastructure (dubbed Katstore) manages the volume, velocity and variety of this data. Katstore is comprised of several stages of data collection and transportation. The stages move the data from monitoring nodes to storage node to permanent storage to offsite storage. Additional information (e.g. type, description, units) about each sensor is stored with the samples.

INTRODUCTION

On each node in the CAM system a monitoring process is responsible for collecting sensor samples for storage. The monitor process communicates with the proxy processes that, in turn, communicate directly with the devices and other systems. The rates at which the monitor processes collect these samples are configured upfront in the central configuration system.

A sample consists of the *sensor name*, *sample timestamp*, *value timestamp*, *status* and *value*.

Sample timestamp is the time at which the CAM system received the sample reading from the sensor. The *value timestamp* is the time at which the acquisition was performed on the sensor and the *value* stored. The *status* field holds the status of the sensor. Timestamps are represented as the time in seconds since the epoch of 1 January 1970 00:00:00 UTC.

The sensor sample storage system is responsible for collecting the samples from the monitor processes. The storage system transports the samples to a central storage node from where they can be queried and archived.

ARCHITECTURE

The storage system is comprised out of several elements (Fig. 1). Each element performs a specific task.

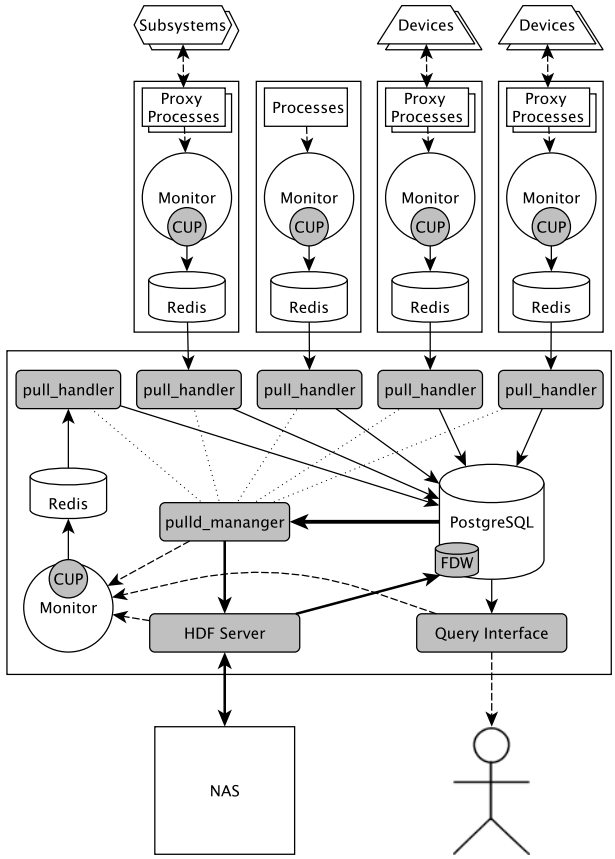


Figure 1: Connections between elements of the storage system.

Memory Buffer

Redis [4], an open source memory database server, is installed on each of the CAM nodes. Redis acts as a buffer to which the monitor process on the node writes sensor samples. On startup and at intervals, the monitor process will write metadata (e.g. type, unit of measure, description etc.) of the sensors to Redis. Internal to the monitor process, a class was developed to manage the writing to the memory database. This cache update class (CUP) runs in its own thread.

Pull Daemon - Pulld

A set of processes on the storage node pulls the samples and metadata out of Redis and stores it into the central database. These processes are collectively called Pulld. On a scheduled basis (currently daily), Pulld analyses the

* martin@ska.ac.za

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

UPGRADES TO THE INFRASTRUCTURE AND MANAGEMENT OF THE OPERATOR WORKSTATIONS AND SERVERS FOR RUN 2 OF THE CERN ACCELERATOR COMPLEX

A. Bland, S.T. Page, CERN, Geneva, Switzerland

Abstract

The Controls Group of the CERN Beams Department provides more than 400 operator workstations in the CERN Control Centre (CCC) and technical buildings of the accelerators, plus 300 servers in the server room (CCR) of the CCC. During the long shutdown of the accelerators that started in February 2013, many upgrades were done to improve this infrastructure in view of the higher-energy LHC run. The Engineering Department improved the electrical supply with fully redundant UPS, on-site diesel generators and for the CCR, water and air cooling systems. The Information Technology Department increased network bandwidth for the servers by a factor of 10 and introduced a pilot multicast service for the video streaming of the accelerator status displays and beam cameras. The Controls Group removed dependencies on network file systems for the operator accounts they manage for the Linacs, Booster, PS, ISOLDE, AD, CTF3, SPS, LHC and cryogenics. It also moved away from system administration based on shell scripts to using modern tools like version-controlled Ansible playbooks, which are now used for installation, day-to-day re-configuration and staged updates during technical stops.

INTRODUCTION

The Controls Group of the Beams Department at CERN manage disk-based Linux workstations and servers for the operation of all of CERN's accelerators, cryogenics and technical infrastructure.

The workstations are business-class Personal Computer (PC) desktops ranging from 2006 to 2015 models running Scientific Linux CERN 6 (SLC6, based on Red Hat Enterprise Linux 6). Eighty of them are concentrated in the CERN Control Centre (CCC) on the Prévessin site in France and the rest are distributed across more than 100 technical buildings including underground areas of the Large Hadron Collider (LHC). All workstations can run the software of any part of the accelerator complex; restrictions on what can be done from where are enforced by the Role-Based Access Control system.

The servers are almost entirely in the CERN Control Centre server room (CCR) with a small number two kilometres away on the Meyrin site in Switzerland for redundancy and recovery purposes. The majority were upgraded to SLC6, only ten still run SLC5.

Although the accelerators themselves have well-defined year-end technical stops and around 2 or 3 technical stops during the year for upgrades, the technical infrastructure operation (electricity, cooling and ventilation and safety installations) and cryogenic parts of the control system must run 24 hours a day, 365 days per year.

TIMELINE

Most of the upgrades were performed between the start of Long Shutdown 1 (LS1) [1] in February 2013 and the first LHC circulating beams of Run 2 in April 2015, see Table 1. Not all of this time was suitable for infrastructure upgrades as the injectors for the LHC, but not the LHC itself, ran from April 2014 to December 2014.

Table 1: Principal Dates

Date	Activity
Q1 2006	First accelerator operation from the CCC
Q3 2008	LHC first circulating beams
Q1 2010	LHC first physics at 7 TeV
Q1 2013	Long Shutdown 1 starts
Q2 2014	Injector startup
Q2 2015	LHC Run 2: first circulating beams and physics at 13 TeV
Q1 2019	Long Shutdown 2 expected, 18 months duration [2]

Another complication for upgrades between the start of LS1 and April 2014 was that not all accelerators stopped; the Compact Linear Collider Test Facility (CTF3) and Linear Accelerator 4 (Linac4) tests had to continue, plus the LHC Magnet Test Bench (MTB) continued operation.

OPERATIONS REDUNDANCY

The Technical Infrastructure (TI) operator environment has been replicated in the fire station on the Meyrin site with four workstations, each with two screens. The main TI telephone number can be switched to this installation. One TI shift per month is run from the fire station to ensure that all activities are possible from there.

Similarly the firemen's environment in the fire station has been replicated in the CCC with five desktops, the Terrestrial Trunked Radio (TETRA) system and the hardwired fire alarm panel. One shift per week is run from the CCC to test the installation.

ELECTRICAL NETWORK

From first operation of the CCC in 2006 up to LS1, powering was already possible either from the Swiss or French national electrical network. Diesel backup if needed came from the Meyrin site and three sets of Uninterruptible Power Supplies (UPS) were locally installed on the Prévessin site providing N+1 redundancy.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

PAST, PRESENT AND FUTURE OF THE ASKAP MONITORING AND CONTROL SYSTEM

Malte Marquarding, CSIRO Astronomy and Space Science, Epping, Australia

Abstract

The Australian Square Kilometre Array Pathfinder (ASKAP) is CSIRO's latest radio interferometer located in the Mid West region of Western Australia. It is in the final phase of construction with Early Science Program starting in the middle of 2016. This fully remotely operated telescope is one of the pathfinder telescopes of the Square Kilometre Array (SKA) building knowledge about infrastructure, technologies and scalability. This paper presents another incremental status update on previous reports. It focuses on software lessons learned and modifications arising from the initial six antenna test system. ASKAP's software consists of two major components, the Telescope Operating System (TOS) and the Central Processor (CP). This paper addresses the TOS, highlighting the monitoring and control aspects of the system.

PROJECT STATUS

The full ASKAP telescope consists of 36 dishes of 12m in diameter hosting phased array feed (PAF) detectors. The hardware is located in a specially designed central building which addresses self-generated radio interference. An initial set of six PAFs, the Boolardy Engineering Test Array (BETA) has been installed and commissioned. The evaluation of this system lead to an improved design with new 36 beam PAFs and an updated digital back end including RF over fibre. This is called ASKAP design enhancement (ADE) and resulted in most of the hardware being moved from the antenna dish pedestals to the central site building connecting to the dishes through fibre optic links. 12 ADE antennas with end-to-end integration of TOS and CP will form the Early Science Program platform. 25% of the total time will be assigned to science project during the time the rest of the telescope will be commissioned. It will produce full data products with only the reduced number of antennas separating it from full ASKAP operation.

Currently the project has evolved to encompass the following active test and production platforms.

- A Parkes 12m dish hosting prototype a PAF. This system is dedicated to engineering tests and new PAF developments such as improvements to the current system as well as next-generation SKA prototyping. Several prototypes have been tested. The software to operate this facility has been stable and remained largely unchanged for several years.
- The BETA system at MRO which host Mark 1 phased array feeds. It is used for beam forming research and small science projects run through the ASKAP Commissioning and Early Science (ACES) team. It has produced several reviewed astrophysics publications and has proven to be an invaluable tool for control system integration, deployment procedures and data management.
- A Mark 2 PAF prototype system at MRO. This single antennas system is dedicated to engineering testing. It is also used for testing hardware modification to the existing Mark 2 set up. Most recently feedback from ACES has prioritised the use of on-dish calibration devices (ODC) to simplify complicated calibration procedures. One such system is being commission on this antenna.
- A four antenna array hosting Mark 2 phased array feeds. It forms the base of the full production array. Commissioning of these systems is taking place at the moment encompassing interferometer phase closure, moving to 36 electronically formed beams and visibility streaming.

For more information about ASKAP visit the project's home page [1].

PROGRESSING THE ASKAP MONITORING AND CONTROL SOFTWARE

Architecture and Framework Update

As described above ASKAP project went through three major stages. A prototype phased array feed on a non ASKAP antenna co-located with the Parkes 64m telescope which still used for research into beam weight generation and stability. Beam weights are used to create beams according to constraints such as maximising the signal-to-noise for a given configuration. This engineering test system was also used to prototype the software technologies and formed the base line of the current system. As described in our previous update [2], the ASKAP software architecture is supported by two key technologies, EPICS and ZeroC ICE. EPICS handles the control and monitoring aspects through software input/output controllers (IOC), whereas ICE defines the interface between software components. Ice interfaces provide the main service contracts between the Central Processor and the Telescope Operating System (see Fig. 1).

We have been using the EPICS framework since the beginning of the project and have now added several helper modules to abstract common usage in ASKAP. For engineering and operator displays of monitoring and control ASKAP has now fully adopted Control Systems Studio (cs-studio). We have joined the project and ASKAP has become one of the products of the upstream source. As part of the collaboration we are in the final stages of setting up cloud-based continuous integration with open access to all collaboration partners. In addition to working with the cs-studio core pack-

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

THE LASER MEGAJOULE FACILITY: CONTROL SYSTEM STATUS REPORT

J. Nicoloso, CEA/DIF, Bruyères le Châtel, 91297, Arpajon, France

Abstract

The Laser MegaJoule (LMJ) is a 176-beam laser facility, located at the CEA CESTA Laboratory near Bordeaux (France). It is designed to deliver about 1.4 MJ of energy to targets, for high energy density physics experiments, including fusion experiments. The commissioning of the facility was achieved in October 2014. This paper gives an overview of the general control system architecture, which is designed around the industrial SCADA PANORAMA, supervising about 500 000 control points, using 250 virtual machines on the high level and hundreds of PCs and PLCs on the low level. We focus on the rules and development guidelines that allowed smooth integration for all the subsystems delivered by a dozen of different contractors. The integration platform and simulation tools designed to integrate the hardware and software outside the LMJ facility are also described. Having such tools gave us the ability of integrating the command control subsystems regardless the coactivity issues encountered on the facility itself.

LMJ FACILITY

The LMJ facility covers a total area of 40,000 m2 (300 m long x 150 m wide). It is divided into four laser bays, each one accommodating 5 to 7 bundles of 8 beams and a target bay holding the target chamber and diagnostics. The four laser bays are 128 m long, and situated in pairs on each side of the target chamber. The target bay is a cylinder of 60 m in diameter and 38 m in height. The target chamber is an aluminium sphere, 10 m in diameter, fitted with several hundred ports dedicated to laser beams injection and diagnostics introduction. Numerous diagnostic instruments are placed in the target chamber around the target to record essential measurements and observe the target behaviour during its implosion. These diagnostics are the prime tools for the physicists to determine the characteristics of the plasma under study.

LMJ CONTROL SYSTEM

LMJ Control System functions

The main functions of the control system are shots execution and machine operations: power conditioning controls, laser settings, laser diagnostics, laser alignment, vacuum control, target alignment, target diagnostics.

All these components are triggered with a high precision Timing and Triggering system [1].

The control system has also a lot of other major functions: personnel safety, shot data processing, maintenance management.

Conducting a shot is composed of two phases: first a master countdown prepares the machine and secondly an automatic sequence [3] executes the shot from the power conditioning charging to the target implosion.

The master countdown has an expected duration of about four hours and the final automatic sequence lasts a few minutes from power conditioning charging to shot execution.

The master countdown coordinates manual operations or automatic programs that prepare the machine: automatic settings computation [4] and associated downloading, laser and target alignment, diagnostics preparation. This can take 2 or 4 hours.

Then, when the laser is ready, the automatic sequence is started: the power conditioning is charged. This takes a few minutes. Then the computer system hands over to the electronic timing system that guides the laser pulses from the master oscillator sources to the target through the amplifiers and transport sections. This takes about 1 microsecond.

General Architecture

The LMJ control system has to manage over 500 000 control points, 150 000 alarms, and several gigabytes of data per shot, with a 2 years on line storage.

It is composed of a dozen of central servers supporting about two hundreds of virtual machines at the central controls level and about 450 PLC's or rack mount PC's at low levels.

Hardware Architecture

From the hardware point of view the LMJ control system is constituted of two platforms located in two different buildings:

- one for system integration (PFI), which is in operation in a dedicated building and consisting of a clone of the operational control system at the supervisory levels and a mixture of simulators and real controllers for representing low levels controls and real equipment [2].
- The operational platform, consisting of two sub-platforms: a small one for integrating the laser bundles (PI) and one for normal operations (PCI).

For each platform, two redundant Alcatel-Lucent OmniSwitch cabinets provide redundant Gigabit attachments to twelve subsystems backbones and main servers (Fig. 2).

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

OVERVIEW AND STATUS OF THE SWISSFEL PROJECT AT PSI

Markus Janousch, Arturo Alarcon, Kristian Ambrosch, Damir Anicic, Alain Bertrand, Kurt Bitterli, Helge Brands, Patric Bucher, Tine Celcer, Pavel Chevtsov, Edwin Divall, Simon Ebner, Martin Gasche, Alexandre Gobbo, Colin Edward Higgs, Fabian Hämmerli, Thomas Hövel, Tadej Humar, Guido Janser, Gaudenz Jud, Babak Kalantari, Rene Kapeller, Renata Krempaska, Daniel Lauk, Michael Laznovsky, Christian Lüscher, Hubert Lutz, Dragutin Maier Manojlovic, Fabian Märki, Trivan Pal, Werner Portmann, Simon Rees, Thierry Zamofing, Christof Zellweger, Dirk Zimoch, Elke Zimoch , Paul Scherrer Institut, Switzerland

Abstract

Recently, the installation of the components for the free electron laser SwissFEL has started at the Paul Scherrer Institute (PSI). In March 2016, beginning of the injector commissioning is planned and first lasing is foreseen a year later. New hardware, like VME64x-boards (IFC 1210, an P2020 based intelligent FPGA controller from IOxOS) and -crates (Trenew), timing system (from MRF with advanced features), motion controllers (Power PMAC from Delta Tau, and MDrive from Schneider), among others, as well as modern field buses, pose great challenges to the controls team. The close interaction of machine- and experiment-components require advanced software concepts for data-acquisition, -distribution, and -archiving. An overview of the project will be presented and the different HW and SW solutions based on the experience gained from preliminary implementations at other facilities of PSI will be explained. First results of the HW commissioning at the SwissFEL will be reported.

INTRODUCTION

The Paul Scherrer Institut (PSI) is the largest research centre for natural and engineering sciences within Switzerland and is located northwest of Zurich, about an half hour's drive away. Research is performed in three main subject areas: Matter and Material; Energy and the Environment; and Human Health. For this fundamental and applied research we use three different accelerator facilities on site and support additional test facilities:

- The facilities consist of one of the highest intensity 570 MeV proton cyclotron, HIPA, that produces particles for fundamental particle physics, muons, and is the source for the neutron spallation source SINQ.
- Another low energy superconducting proton cyclotron is used for the patient treatment systems at the Center for Proton Therapy.
- The Swiss Light Source, SLS, is a 3rd generation synchrotron and houses 18 beamlines.

These accelerators are available for the international community; we host about 2400 external users per year. All these facilities are controlled from a central control room.

Until the end of October 2014 a 250 MeV electron Linac, SITF (SwissFEL Injector Test Facility), was in operation and was used as a pre-study for the free electron laser (FEL) project. Many of its components are used as the injector for the upcoming FEL. The new SwissFEL is currently set-up a few hundred meters east of the main campus of PSI.

Of course, there are several other test-facilities and laboratories related to the accelerator facilities.

This large array of different facilities and set-ups around the laboratory poses quite a challenge to be handled by one Controls group. Therefore a lot of emphasis has been put on standardization on the equipment that is supported and the tools that are being used within the Controls group. However, there are new experiments and test set-ups constantly coming up that need to be supported on the one hand, on the other one has to watch a growing heap of legacy components and code that have to be maintained. E.g. HIPA is operational since over 40 years.

Standardization in the Past

In order to minimize any overhead and maximize the efficiency of the Controls group a strong standardization effort is implemented. Among these standards are: All accelerator and most control systems are based on EPICS [1]. For fast I/O, scalers, etc. we use a VME bus architecture, controlled by a Motorola CPU board running VxWorks 5. The timing system is based on the model 230 of Micro-Research Finland [2]. The main motion controller used is MaxV from Pro-Dex.

Consoles and most servers run Scientific Linux 6.4 and the Linux software is distributed through Puppet. To deploy software for the EPICS system a mechanism developed at PSI is used. A relational database keeps an overview of our hardware inventory and helps in managing the many components at the facilities.

THE SWISSFEL PROJECT

The new SwissFEL is currently set-up a few hundred meters east of the main campus of PSI (Fig 1). The project consists of two consecutive phases.:

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

COMMISSIONING OF THE TPS CONTROL SYSTEM

C.Y. Liao#, J. Chen, Y.S. Cheng, P.C. Chiu, C.Y. Wu, C.H. Huang, K.H. Hu, Y.T. Chang,
D. Lee, S.Y. Hsu, C.J. Wang, C.H. Kuo, K.T. Hsu
National Synchrotron Radiation Research Center, Hsinchu 30076, Taiwan

Abstract

Control system for the Taiwan Photon Source (TPS) has been completed in 2014. Commissioning of the accelerator system is in proceeding. Electron beam were stored at the storage ring and emit first light in December 31, 2014. TPS control system adopts EPICS toolkits as its frameworks. The subsystems control interfaces include event based timing system, Ethernet based power supply control, corrector power supply control, PLC-based pulse magnet power supply control and machine protection system, insertion devices motion control system, various diagnostics related control environment, and etc. The standard hardware components had been installed and integrated, and the various IOCs (Input Output Controller) had been implemented as various subsystems control platforms. Low level and high level hardware and software are tested intensively in 2014 and final revise to prepare for routine operation is under way. Efforts will be summarized at this paper.

INTRODUCTION

The TPS [1] is a latest generation of high brightness synchrotron light source which is constructed at the National Synchrotron Radiation Research Center (NSRRC) in Taiwan. TPS consists of a 150 MeV electron linac, a booster synchrotron, a 3 GeV storage ring, and experimental beam lines. Ground breaking for civil construction was held on February 2010. The construction works were completed in April 2013. Accelerator system installation and integration was proceeding in later 2013. The control system environment was ready in half of 2014 to support subsystem integration test and commissioning without beam. After 4 months of hardware testing and improvement, the TPS initiated the commissioning of the booster ring and storage ring on December 2014. On the last day of 2014, the TPS has delivered its first synchrotron light [2, 3].

Adequate and reliable functionality of control system is one of the key to the success of TPS commissioning. Control system for the TPS is based on the EPICS framework [4]. The EPICS toolkits provide standard tools for display creation, archiving data, alarm handling and etc. The EPICS is based on the definition of a standard IOC structure with an extensive library of driver and support a wide variety of I/O cards. The EPICS toolkits have various functionalities which are employed to monitor and to control accelerator system.

The TPS control system consists of more than a hundred of EPICS IOCs. The CompactPCI (cPCI) is equipped with input/output modules to control subsystems as standard IOC. The other kinds of IOCs are also supported by the TPS control system, such as BPM IOC, PLC IOC, various soft-IOC and etc.

To achieve high availability of the control system, emphasis has been put on software engineering and relational database for system configurations. Data channels in the order of 10^5 will be serviced by the control system. Accessibility of all machine parameters through control system in a consistent and easy manner will contribute to the fast and successful commissioning of the machine. High reliability and availability of TPS control system with reasonable cost and performance are expected.

SYSTEM COMPONENTS

The system installation and integration with subsystems for the control system was done [5]. Details of the control system are summarized in the following paragraph.

General EIPCS IOC Interface

There are many different kinds of IOCs at equipment layer to satisfy various functionality requirements, convenience and cost consideration, shown in Table 1. Most of the devices and equipments are directly connected to cPCI IOCs with EPICS. The cPCI EPICS IOC is equipped with the cPCI-6510 CPU board. The cPCI-7452 128 bits DI/DO module is used for BI, BO solution. ADC and DAC modules in IP (Industry pack) module form-factor are used for smaller channel count application, such as insertion devices control. Event system modules are in 6U cPCI form-factor. Private Ethernet will be heavily used as field-bus to connect many devices. Power supplies of all magnets except for correctors are equipped with Ethernet to the EPICS IOC. Multi-axis motion controller with Ethernet interface is the standard for the control system.

Ethernet attached devices are connected to the EPICS IOC via private Ethernet. Devices support VXI-11, LXI, Raw ASCII and Modbus/TCP protocol are connect to EPICS IOC directly by TCP/IP interface. Devices of this category include power supply, temperature acquisition (RTD or thermocouple), digital multi-meters, oscilloscopes, signal generator, and other instruments.

All corrector power supplies are driven by the corrector power supply controller (CPSC) module [6]. The CPSC equips with 20 bits DAC and 24 bits ADC. Two SFP ports supported by the on board FPGA (Spatan 6), these SFP ports are receive correction setting (Aurora and Gigabit Ethernet by using UDP/IP protocol) from fast orbit feedback FPGAs to slow orbit feedback PC, feed-

#liao.cy@nsrrc.org.tw

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

STATUS OF THE EUROPEAN SPALLATION SOURCE CONTROL SYSTEM

T.Korhonen, R. Andersson, F.Bellorini, S.Birch, H.Carling, J. Cereijo García, L. Fernandez, R. Fernandez, B.Gallese, S. Gysin, E. Laface, N.Levchenko, M.Mansouri, A. Monera-Martinez, R. Mudingay, A. Nordt, D. Paulic, D. Piso, K. Rathsman, M. Rescic, G. Trahern, M. Zaera-Sanz
European Spallation Source, Lund, Sweden
K. Strnisa, U.Rojec, N.Claesson, A.Söderqvist, Cosylab, Ljubljana, Slovenia

Abstract

The European Spallation Source (ESS) is a collaboration of 17 European countries to build the world's most powerful neutron source for research [1]. ESS has entered the construction phase and the plan is to produce first neutrons by the year 2019 and to complete the construction with 22 neutron instruments and reach the target 5 MW beam power by 2025. The Integrated Control System Division (ICS) is responsible to provide control systems for the whole facility.
The unprecedented beam power and the construction of the facility as a collaboration of the member countries, with many components delivered in-kind presents a number of challenges to the control system. The Integrated Control System Division has to prepare systems and specifications so that the work can be effectively shared between the contributors and staff on-site. The systems need to provide sufficient performance for successful operation of the facility, be standardized to a level that allows easy integration into the facility during a short installation period and have to be maintainable by the in-house staff after the construction has finished.

INTRODUCTION

The European Spallation Source (Fig. 1) is an accelerator-based facility currently in construction in Lund, Sweden. The facility will provide neutron beams for various fields of research, using a variety of instruments in neutron beam lines. Different components of the ESS will be delivered in-kind by several institutions in the member countries. This applies to all parts of the project, including the control system.

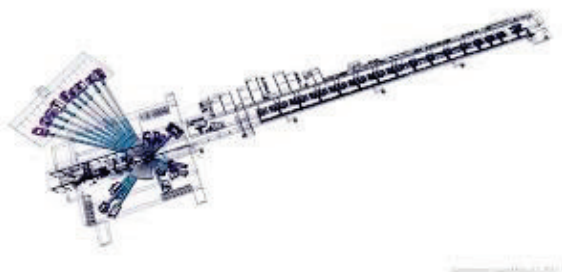


Figure 1: A schematic view of the ESS facility.

In ESS, a high power proton beam hits a spallation target to produce neutrons. The proton beam is generated in a proton linac (Fig. 2) which operates at 14 Hz repetition rate and 5 megawatt average (125 MW peak) beam power. The about 600 meters long linac consists of a proton source, normal conducting ("warm") section, a superconducting section and finally a high energy beam transport section towards the target [2].

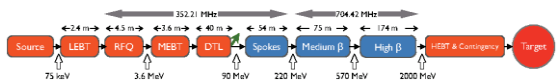


Figure 2: The ESS linear accelerator.

The ESS accelerator produces long pulses of 2.86 milliseconds in normal operation. This pulse lenght allows a great flexibility for designing the neutron experiments.

The accelerated proton beam hits the neutron target (Fig. 3), a rotating tungsten wheel divided into several sectors. The produced spallation neutrons go through moderators and neutron guides to the neutron instruments. The accelerator operation has to be synchronised with the rotation of the wheel so that the beam pulses hit the middle of each sector. A rastering magnet system expands the beam before it hits the target so that the heat load is distributed to a wider area.

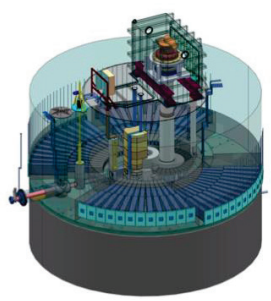


Figure 3: The neutron target of ESS.

Neutron beam lines carry the produced neutrons to the instruments for detection. The neutron beam lines have among other components several choppers on the way to select the neutron energy, or wavelength. The choppers have also to be synchronised with the beam operation.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors