IEEE Transactions on Nuclear Science, Vol. NS-28, No. 3, June 1981

THE DEVELOPMENT OF AN INTELLIGENT RELIABLE DATABASE DRIVEN ALARM SYSTEM AT THE ISR

C. Fischer, J.N. Gamble
European Organization for Nuclear Research (CERN)
1211 Geneva 23, Switzerland

Summary

The monitoring of equipment and operating conditions of a complex installation is an important task for an accelerator control system, especially for a machine with high intensity circulating beams like the ISR. It is essential that the control system provides the relevant correct alarms and warnings whilst at the same time not flooding the operators with irrelevant information.

At the ISR a system has been introduced which displays fault conditions as text messages on a colour display. In conjunction with these messages it is possible for the operator to obtain supplementary information concerning the fault or actions that should be taken. The operator also has facilities to cancel fault messages, to ignore them and to supervise the monitoring programs.

The texts for the messages are kept on a database in order to facilitate their modification. All fault messages are printed and kept on a file for later extraction by an interrogation program. This allows for the easy analysis of fault situations or for the production of statistics.

Introduction

The need to have a good alarm system that is reliable, gives concise accurate information on faults, and also provides information on the remedial actions to take has become very important for the ISR. Some aspects of storage ring operation have strongly influenced the design of the alarm system described here. Proton beams of more than 35 Amps are used at the maximum energy of 31.4 GeV for experimental physics. These high intensities are usually achieved after periods of preparation of the order of ten to twelve hours[1]. In such a time scale the tracking of faults is of vital importance in terms of operation time saving. For the ISR one can say that this is the major function of the alarm system as the safety of the machine is generally ensured by proper hard-wired interlock chains on the equipment.

An alarm system reflects the equipment of the accelerator and as such it is always in a state of change, new equipment is installed and old indications can take on new meanings. This continual change puts a load on the operators to be able to respond to alarms in the correct way. As greater and greater demands are made on the performance of the accelerator so the number of monitored quantities increases. With already several hundred different types of alarm condition, referring to many more items of equipment, the task of conveying the urgency and the root cause of an alarm in a simple text message is of considerable importance. In addition enough flexibility must be provided to implement any required modifications without a major rework of the software.

These problems have been tackled at the ISR and this paper is intended to give an insight into the alarm system and the facilities it offers.

The previous alarm system

There has been an alarm system in the ISR since the beginning of the exploitation of the machine. In 1971 the monitoring of the vacuum system was introduced as one of the first group of monitoring programs that make up the alarm system. All the alarm conditions were reported to a central printer. Monitoring programs have since developed regularly through the years either as new equipment was introduced, or as the existing equipment became more critical in the ever increasing performance requirements. Around 1973 they were performing a limited consequential analysis on the state of the equipment and modifying the type and priority of alarms according to the operational phase of the machine (i.e. Machine development, setting up, stable beams). Each program was (and still is) activated via switches under the control of the operator, giving a wide range of flexibility in the monitoring that should be done both during machine running periods and during controlled accesses or shut downs.

The incentive to develop a new alarm system came for several reasons, mainly related to the fact that the output device for all the alarms was a printer. Furthermore this situation was considerably worsened by the lack of guidelines and standards for the presentation of the alarm messages. This resulted in the following drawbacks :

- Too many messages were output corresponding to either false alarms, redundant information or information of use only by specialists. A real fault could easily be masked and the system was not taken seriously enough.

- The text of the messages was presented differently by each monitoring program making the analysis of the printer output difficult. This could hamper the detection of a particular fault and made post mortem investigations difficult. In spite of the fact that the monitoring programs issued "O.K." messages as faults were corrected, it remained extremely difficult to keep track of the outstanding alarms on the printed output.

- No possibility was offered to suppress recurrent false alarms other than by removing the monitoring program and thus preventing a complete system from being monitored.

- Following the modification of a system or the installation of a new one, it was often a non-trivial job to implement the necessary arrangements.

With these problems in mind work was started on the new alarm system towards the end of 1978 when the main requirements were discussed. The actual alarm system was introduced at the end of 1979 and by the start of 1980 most of the monitoring programs had been converted to the new system.

The new alarm system

The alarm system is naturally an integral part of the ISR control system [2,3] and has evolved alongside the more recent developments. Consequently many of the tools that directly or indirectly have contributed to the alarm system were developed for the control system as a whole. The database software, display software and message transport system were readily available when the design of the alarm system was made. The major

0018-9499/81/0600-2341$00.75©1981 IEEE

consequence of this was that the design could concentrate on the operational requirements, without becoming overwhelmed by problems of implementation.

The alarm system is made up of three levels or types of building blocks. At the top there is the operator interface dealing with the presentation of the messages and the actions that the operator requests. Here all the alarms messages are collected together for processing and display. The arriving messages are already in plain text accompanied by some control information. Of the two other levels, one is concerned with the detection of the alarm conditions and with consequential and reduction analysis, and the other is responsible for the production of the textual messages.

The two areas of importance are felt to be, firstly the operator interface, not so much for the hardware involved but for the facilities offered, and secondly the use of databases in the analysis of faults and the production of text messages.

The operator interface

The ISR is operated from two main consoles and each of these is equipped with a colour screen dedicated to faults, this screen being the interface between the alarm system and the technician operating the machine. The information about a fault is therefore seen immediately and an audio alarm is emitted to draw the attention of the operator. For the display of the alarm messages both screens give exactly the same information. Below each is situated a row of buttons through which the operator exercises control over the display and the processing of alarm messages. In addition a printer is used to record the arrival of all fault messages and their corresponding resets. More detailed information can also be provided by the monitoring programs for the use of the hardware specialists which would be too detailed for the operations crew. This auxiliary information is kept by the alarm system and can be displayed on the screen if required.

Faults may arise for a variety of reasons that can be classified as either occurring during an action requested by an operator or as happening independently. The philosophy adopted in the alarm system was that only independent asynchronous faults should be dealt with. Problems encountered during the execution of an application program are relayed immediately to the operator concerned by the program, and do not appear as faults on the alarm screens. The asynchronous faults are for operational convenience classified into two categories :

- ALARM a major fault requiring immediate intervention of the operator to preserve the operating state of the machine.

- WARNING a situation that while not immediately serious could develop into a major fault if not corrected.

All of the messages are presented by the alarm system in a standard format each occupying one line of the display. This presents all common information, time, program name, subsystem affected at the beginning of the text message and is suitable, not only for a quick appraisal by the operator, but also as it improves immensely the analysis of the printed output.

An important requirement of an alarm system is that the operator should be able to understand quickly the nature of a fault. Two steps that progress towards this goal are to arrange that the messages are concise and accurate, and secondly to avoid flooding the operator with messages. As an additional requirement it was decided to insist that as far as possible all the outstanding alarms should be accommodated on a single screen. This may seem to require an intelligent approach to reduce the number of messages, but was in fact achieved by using the simple method mentioned below.

At the lower levels of the alarm system the equipment is monitored by individual programs which deal with complete subsystems. Taking advantage of this, the alarm processor limits the number of messages on the screen from a particular monitoring program. In the present system a maximum of three are shown, and if more are present then a summary message is shown instead. All the details associated with a summary message may be retrieved at the touch of a button.

In most cases when a fault disappears the monitoring program which generated the original alarm informs the alarm system, and the corresponding message is then erased. However some types of fault message cannot be cancelled automatically (for example partial beam current loss) as it is difficult to define what the correct conditions is. This is typical for situations where changes are monitored rather than reference values. Facilities are therefore supplied to allow these messages to be manually cancelled.

When, in spite of all precautions and filtering mentioned, a fault condition is continually signalled and then reset without affecting the operation of the machine (as can be the case with a faulty acquisition system or intermittent changing of status bits in the hardware) then the alarm system can be instructed to ignore that fault. Later, when the hardware has been repaired, the fault may be again activated. The number of such ignored faults is permanently displayed on the screen and the operator can see the full list at the touch of a button.

The HELP facility. When the present alarm system was built, an operational procedure was established for each individual fault giving briefly the meaning of the message, and also the action that should be undertaken when it appears. These procedures were classified in a folder which could be referred to when needed. The HELP option, implemented at the end of 1979 fully computerised this facility. For each message an operational procedure is catalogued in the alarm system database and can be made immediately available on the screen.

Databases and the Low Level Alarm System

One of the major problems tackled was that of program modification due either to hardware changes being made, or to an increase in the number of items of equipment monitored, or simply to changes in the textual content of messages. This problem is compounded by the fact that consequential and reduction analysis may be performed by the monitoring program. A common solution using databases was adopted for both.

The ISR accelerator is notationally divided into octants and further into sectors. Taking for example the requirements for vacuum monitoring, alarm messages should be given when an individual pump fails, but if all the pumps in a sector fail then only a relevant summary message should be given.

Similarly for the magnet power supply monitoring system global effects of interlock chains have to be considered. If a water circuit fails then all the power supplies on that circuit must switch off. Fault messages have to be generated for the supplies that remain on, whereas normally it is an alarm for a power supply that switches off.

Over the whole of the ISR this interplay of consequences and reduction posed a non-trivial problem for the alarm system. The solution adopted was to use database software to fully define the elements of the accelerator and their interrelationships. However the application of this information was left to the individual programs that were intrinsically involved with the various types of equipment that they monitored. The databases operated at two levels. The lower level database, which already existed for the control system, described the equipment giving its address, units of measurement, scaling factors etc. and several control parameters. This information was expanded to include such items as the water circuit, octant, sector etc. needed for the consequential and reduction analysis. The most important items in this database from the alarm system point of view, were the equipment name, its address and the information used in performing the reduction and consequential analysis. As the central alarm system dealing with the displays also enforces a global reduction in the number of displayed messages and performs a limited consequential analysis, the monitor program writer is allowed to decide whether it is justified to add sophisticated consequential analysis, or if it is sufficient to use the default facilities. At the present time about half the monitoring programs incorporate their own consequential and reduction analysis.

In addition to the equipment database a new database was introduced exclusively for the alarm system. This contains the text of the alarm messages. Clearly the provision of a unique message for every fault on every piece of equipment is avoided. Each message in the database represents a format rather like the FORTRAN FORMAT. The monitoring program supplies the format number and a list of parameters to substitute into it. The range of possible representations of the parameters is similar to that of FORTRAN, with the addition of the ability to specify other formats to be imbedded depending on an index given as a parameter. The substitution is performed by the second level of the alarm system. The monitoring program sends the format number and a list of parameters to the local alarm processor which forms the second level. This local conversion keeps the monitoring programs small and enables alarm messages to be recorded on disc, or printed at a local printer on each of the subsystem computers. The correct functioning of each subsystem's monitoring is therefore assured even in times of machine shut down, when cabling and other serious upheavals take place affecting the computer network. This is of particular importance for the ISR vacuum system which is monitored 365 days of the year.

A similar scheme of formats with wild card matching options was adopted to match the text of an alarm message with the index to the operational procedures in the HELP facility. The text is made more meaningful by allowing the wild card matches to be substituted into the text. For example, the alarm message :

2CR108 SWITCHED OFF

may appear for any of the 350 or so power supplies. 2CR108 is one particular power supply. Only one alarm format and one help format will exist for this type of fault. The operational procedure when displayed would, with the substituted wild card text, read :

This means that power supply 2CR108
has switched off ... etc.

Reliability

The lowest level of the alarm system is represented by the many monitoring programs. These run repetitively at various periods ranging from once every 10 to 15 seconds up to once every 10 to 15 minutes. The most common repetition rate being once a minute.

One of the problems inherent in the monitoring of equipment is that when there are no alarm messages being produced one has to be sure that the equipment is still being monitored. With passive subsystems such as vacuum or cryogenics the accelerator can continue to operate even when the subsystem computer has stopped. With the power supply subsystem a computer stoppage would soon be detected but again the control of the accelerator would not be affected by a stoppage of the monitoring programs.

To resolve this question of reliability a surveillance scheme was introduced at the top level of the alarm system. The central alarm system maintains a list of the active monitoring programs and each of them must send a message to indicate that it has finished a check. Overdue conditions can therefore be detected and an alarm message produced. The operator can request a display of the list of alarm programs telling him when they last ran and how frequently they run. The central alarm system itself is surveyed by means of a hardware watch-dog timer, hard-wired to an audio alarm in the control room.

Conclusions

The operations team have now had more than a year's experience with the new alarm system and the expectations have been achieved with a high degree of satisfaction. When important breakdowns affect the machine the filtering of information works very well : the screen display has never overflowed and appears quite sufficient to communicate at any instant all messages concerning major faults. One of the greatest successes was undoubtedly the HELP facility which provided instant instructions on the actions to take when faced with a fault.

Acknowledgements

We should like to acknowledge the previous unpublished analysis and propositions made by J-P. Koutchouk and also the support of R. Keyser. Many other people, from all the groups of the ISR, have contributed to the success of the project, especially in explaining the required fault analysis for their equipment and in the writing of the monitoring programs. The procedures used in the HELP facility were catalogued by J-M. Geroudet.

References

1. C. Fischer et al. : "Performance of the CERN ISR at 31.4 GeV", IEEE Transactions on Nuclear Science, NS-26, N°3, June 1979.

2. J. Gamble et al. : "The extension of the ISR computer control and monitoring system by two NORD-10 computers", CERN-ISR-CO/78-24, October 1978.

3. J. Gamble et al. : "Towards full automation of accelerators through computer control", XIth International Conference on High Energy Accelerators, July 1980.