ISBN: 978-3-95450-262-2 ISSN: 2673-5350 doi: 10.18429/JACoW-IBIC2025-MODC02

PERFORMANCE ANALYSIS OF THE LHC BSRL AND POSSIBLE IMPROVEMENTS *

A. N. Jury†,1,2, A. Hill[1], C. P. Welsch[1], H. D. Zhang[1], D. Butti[2], M. Gonzalez-Berges[2], M. Martin Nieto[2], S. Mazzoni[2], A. Schloegelhofer[2]
[1]University of Liverpool, Liverpool, United Kingdom
[2]European Organization for Nuclear Research (CERN), Geneva, Switzerland

Abstract

The Beam Synchrotron Radiation Longitudinal density monitor (BSRL) at the LHC leverages time-correlated single-photon counting (TCSPC) to provide high-dynamic-range measurements of the relative charges in each RF bucket with a time resolution of 50 ps. These measurements are needed for the operation of the LHC as well as for the luminosity calibration required by the LHC Experiments. In this work, we identify sources of error for each of the BSRL components. These components are the optics (mirrors, filters and optical fibres), the detector (a hybrid photomultiplier - HPM), the electronics (a Time to Digital Converter - TDC), and some data analysis used for the final results. Knowledge of the errors of the BSRL is crucially important as any errors are passed directly into the luminosity calibration of the LHC experiments. We quantify the errors introduced by each of these parts and for external systems, like the LHC timing. For the largest contributors to the overall error, we propose mitigation strategies that can be deployed in the short term.

INTRODUCTION

Not every RF bucket in the Large Hadron Collider (LHC) is intended to be filled. The RF at the LHC's operating energy has a period of 2.5 ns, however, the minimum distance between filled buckets is 25 ns. For most purposes only the filled buckets are of interest. We introduce the concept of an RF slot, a set of 10 buckets arranged as can be seen in Fig. 1. Charges in the nine nominally empty buckets surrounding a filled bucket are referred to as "Satellites". In addition, many slots are also left unfilled. Charges within these unfilled slots are referred to as "Ghosts". Satellites are of concern to the experiments as they contribute to the luminosity background [1].

The BSRL monitors the relative longitudinal charge distributions, including ghosts and satellites, for the identification of operational issues and allows for a correction for satellites in the luminosity calibration. There are two BSRL instruments, one for each beam. The Synchrotron Radiation (SR) used by the BSRL is produced as a beam passes through a bending magnet or a purpose-built superconducting undulator. The latter is necessary to generate visible SR at injection energies [2]. The SR is passed through a range of neutral density filters to control the intensity, followed by a band pass filter to limit the effect of dispersion on the arrival time resolution. The light is then extracted from the tunnel through fibre-optic cables before being detected by a Hybrid Photo-Multiplier (HPM), which is capable of detecting individual photons.Photon counts are given timestamps relative to the start of the turn by the Time-to-Digital Converter (TDC), and histogrammed over many turns to increase statistics, see Fig. 2. From this, the relative populations within each RF bucket can be inferred.

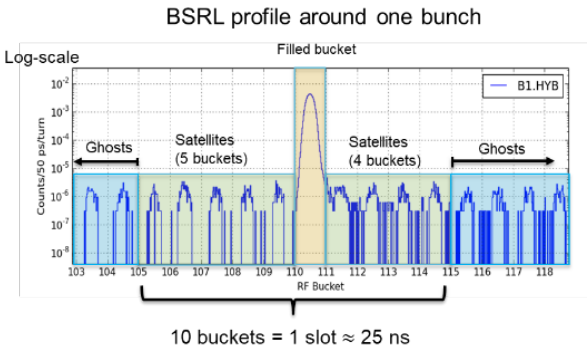

Figure 1: Example of one RF slot, highlighting the definitions of filled, ghost and satellite buckets.

The histogram is built over several minutes. The longer the integration time the larger the dynamic range in counts, and therefore the charges in satellite and ghost buckets are more distinguishable from the background.

COUNTING ERRORS

Due to physical limitations of the detector and the time-to-digital converter, not every photon present at the detector will result in a count, and therefore, true counts will be lost. Equally, due to phenomena such as dark counting and after-pulsing, it is possible to have counts where there is no photon.

Dead Time

After the arrival of a photon, the detector system, composed of the HPM and TDC, will be unavailable to detect any further photons for a given time.

The availability function $S(t)$ is defined as the probability that the detector system is available at some time t after an initial detection.

Left uncorrected, this would cause a distortion in the histogram, favouring counts from the beginning of buckets. As the intensity of filled buckets is much larger than that of

* This work was supported by CERN doctoral student program and Science and Technology Facilities Council under grant agreement ST/W006766/1.
† alexander.jury@liverpool.ac.uk

ISBN: 978-3-95450-262-2 ISSN: 2673-5350 doi: 10.18429/JACoW-IBIC2025-MODC02

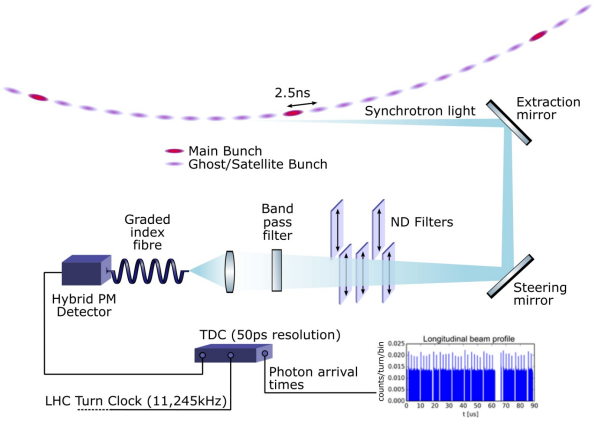

Figure 2: Diagrammatic overview of the BSRL.

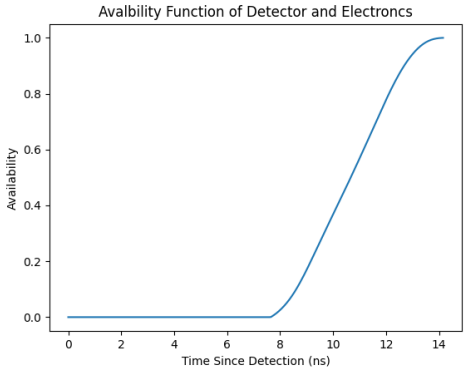

Figure 3: Availability against time after detection of the Acquiris TC890 TDC and PicoQuant PMA06 Detector.

satellites, photon counts from trailing satellites would be suppressed. Thus, it is vitally important that the count rate per bunch per turn is less than one to allow the satellites to be detected. This limited count rate leads to longer integration times and higher statistical errors for the BSRL.

Even with a suitably low count rate, some distortion will still be present. If no part of the signal is completely suppressed, i.e., the availability is not zero on every turn anywhere in the histogram, then with knowledge of $S(t)$ it is possible to remove the effect of dead time from the histogram by applying this formula:

$$\lambda_i = -\ln\left(1 - \frac{c_i}{1 - c_i \otimes (1 - S_i)}\right), \qquad (1)$$

where λ_i is the true mean number of counts per turn in bin i with no dead time, c_i are the measured counts per turn measured in bin i with dead time and $\otimes$ represents discrete convolution. A complete derivation of Eq. (1) can be found in [3].

The availability function for the current BSRL system can be seen in Fig. 3 where the system recovers over the order of 14–15 ns. This comes entirely from the TDC dead time [4] completely masking the dead time of the HPM.

We measure S experimentally via the pump and probe method, with our setup shown in the left panel of Fig. 4.

Using a pair of pulsed diode lasers with a variable time delay τ, the first pulse puts the system into dead time. The second will be detected only at a fraction of all cycles, proportional to the availability at the time delay τ. In this setup, τ is generated electronically. Both signals are then split, with one half subject to a fixed delay greater than the dead time. Meaning that each laser fires twice in each cycle to provide a reference unaffected by dead time.

By comparing the relative number of counts in the probe h and reference probe h_r it is possible to estimate $S(\tau)$. This assumes that the pump is detected on every cycle, so a correction for the efficiency of the probe (total number of counts measured in the probe pulse divided by number of cycles) m, then S as a function of τ is given by Eq. (2).

$$S(\tau) = \frac{1}{m}\left(\frac{h}{h_r} - (1 - m)\right). \qquad (2)$$

In an effort to improve the performance of the BSRL, a new Multiharp TDC from PicoQuant has been purchased. Preliminary investigation using the pump probe method suggests that the dead time of this new system is on the order of 1 ns, with a precision of 5 ps [5]. These facts combined may enable an improved counting rate.

After-Pulsing

An after-pulse is a false count by the detector system, triggered by a true count sometime before, for which there are several possible sources. The HPM consists of a traditional photo-multiplier tube and an avalanche diode (AD). It is possible for charge carriers to become trapped on the wrong side of the junction. These can be freed by thermal excitation at some random time later, and cause a new avalanche. Another source of after-pulsing is back-scattering. Electrons striking the AD can generate secondary electrons. Those electrons are reaccelerated towards the AD by the external voltage to generate false counts. It is also possible that, in the bombardment, an electron strikes a residual gas molecule producing a slowly drifting ion that arrives at the cathode sometime later.

The after-pulsing from the HPM has a distinctive shape as can be seen in Figs. 4 and 5. After-pulsing in the LHC is usually masked by the presence of filled and satellite buckets, but can be observed if the gap between slots is larger than the after-pulsing period, such as in Van der Meer scans, where individual bunches are used.

Figure 5 compares LHC measurements with controlled after-pulsing studies in the lab. The orange curve shows the characteristic after-pulsing signature of the HPM, while the blue histograms show the LHC data. The top panel exhibits only after-pulsing with no additional satellite or ghost structure. In contrast, the bottom panel shows clear ghost bunches between 20 ns and 60 ns. These are distinct from after-pulses because they are well-centred within RF buckets, unlike the broader after-pulse distribution.

ISBN: 978-3-95450-262-2 ISSN: 2673-5350 doi: 10.18429/JACoW-IBIC2025-MODC02

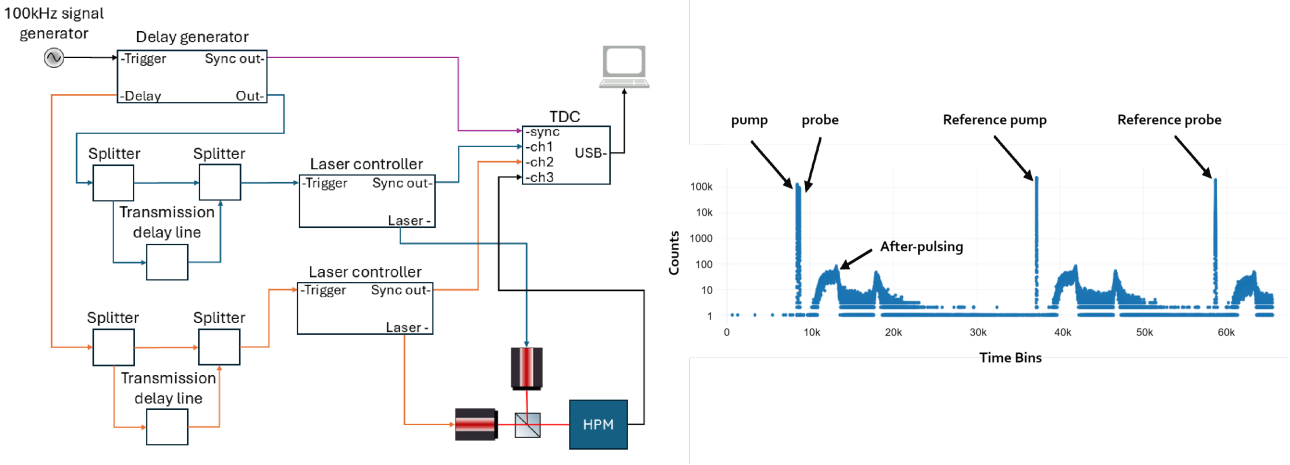

Figure 4: Left: Setup to measure availability function; Right: Example of resultant histogram for a single value of τ.

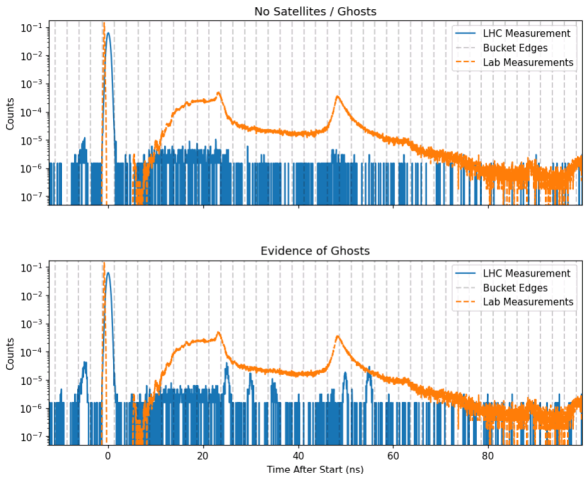

Figure 5: Comparison of BSRL histograms from LHC fill 10826 (blue) and lab after-pulsing measurements (orange). Top: no evidence of ghost bunches; Bottom: ghost bunches visible between 20–60 ns.

Background

Thermal and electrical noise in the system means that even in the absence of light, the detector will report counts. These "dark counts" are primarily thermal in origin [6]. To reduce dark counts, the HPM uses active temperature stabilisation. To correct the histogram for dark counting and after-pulsing, an estimate of the background is made by considering the counts within the separatrix of the local RF buckets, where no particles can reside for many turns, so the majority of counts are due to background. This method of correction can be negatively impacted by any timing errors, causing true counts to be detected in the bucket edge. [3].

TIMING ERRORS

The labelling of counts in time is integral to the creation of the raw histogram. Timing errors come from multiple components, mainly the timing source, the TDC internal clock jitter and the dispersion within the optical fibre. Each

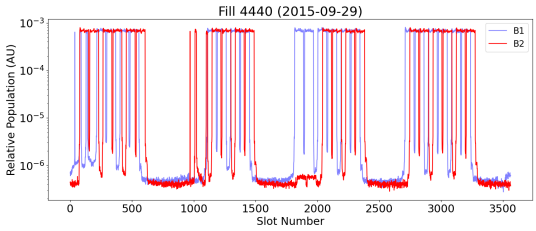

Figure 6: Slot Histogram from LHC fill 4440. Signals due to beam loss can be seen in the region around slot 1000, where the shape of Beam 1 can be seen in Beam 2 [7].

of these sources can be treated independently and so the total timing error is the quadratic sum of all sources.

Optical Dispersion

The optical fibre allows for the extraction of SR from the LHC tunnel to an enclosure within a rack where the HPM is located. This is necessary to protect the detector from particles lost from the LHC. Before 2016 [7], the HPM was located in the tunnel. An echo of the counter-rotating beam could be seen in histograms being produced, such as is shown in Fig. 6.

SR contains a wide range of wavelengths [8]. When the SR is coupled into the optical fibre, the spectral spread leads to chromatic dispersion, causing photon arrival times to vary by wavelength and resulting in spillover of counts between adjacent RF buckets. To minimize this effect, a bandpass filter centred at 532 nm $\pm$10 nm (FWHM) is used to restrict the wavelength range and reduce temporal broadening.

The exact length of the optical fibres for each beam was measured using Optical Time Domain Reflectometry (OTDR), Fig. 7. An estimated value of dispersion within the optical fibre was made by measuring the FWHM of a pulsed diode laser with an HPM and a TDC, with and without the presence of a coupled 15 m patch-fibre of the same fibre type used in the BSRL.

A pulse broadening of 5.182 ps m^{-1} was observed, which was used to estimate the pulse broadening for the full length

ISBN: 978-3-95450-262-2 ISSN: 2673-5350 doi: 10.18429/JACoW-IBIC2025-MODC02

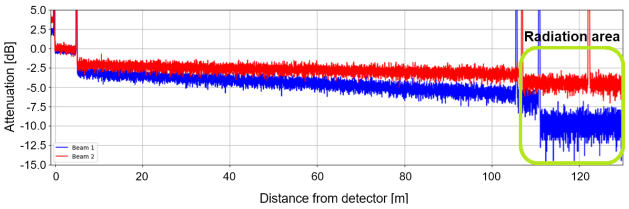

Figure 7: OTDR measurement of the attenuation along the transmission line at a wavelength of 1310 nm.

Table 1: Dispersion Estimates for Beams B1 and B2

Beam	Length	Material Disp.	Total Disp.
B1	110.9 m	484.6 ps	569.1 ps
B2	122.5 m	535.3 ps	628.6 ps

of fibre from each beam to the BSRL, which is the Total Dispersion in Table 1.

The time estimate of pulse broadening due to Material Dispersion as shown in Table 1, was estimated by modelling the fibre as pure fused silica and, considering the wavelength range set by the bandpass filter, is narrowly centred around 532 nm, yielding a value of 437 ps/(nm km) [9].

It is intended to use Second Harmonic Generation Frequency-Resolved Optical Gating (SHG FROG) to find the exact form of dispersion. There are two methods for introducing dispersion compensation under consideration:

- **Physical dispersion compensation**: The uncertainty in the wavelength of a photon present at the detector is translated to an uncertainty in production time. Photons of different wavelengths will take different times to traverse the fibre. This could be remedied by the addition of an equal and opposite dispersion, such that all wavelengths have the same time of flight through our system. This could be achieved through the use of a series of diffraction gratings as shown in Fig. 8
- **Deconvolution dispersion compensation**: An alternative could be to treat the measured histogram as the convolution of the true signal and the dispersion. It may then be possible to restore the signal by deconvolution. This technique is used in Time-resolved fluorescence spectroscopy [10] another field using TCSPC.

Timing Source

All timestamps produced by the BSRL are given relative to a reference marking the start of the turn. Historically, this signal was taken from the Beam Synchronous Timing (BST) [11], which transmits timing information from the RF system to points all around the LHC through kilometres of optical fibre. The uncertainty in this timing is around 1 ns peak-to-peak, or standard deviation of 300 ps, which has been obtained empirically. This error was found to be the dominant timing error in the BSRL. The bunch lengths inferred from the BSRL histogram, compared to those measured by the Beam Quality Monitor (BQM), had a Gaussian blurring applied. The best fit is found when the sigma of

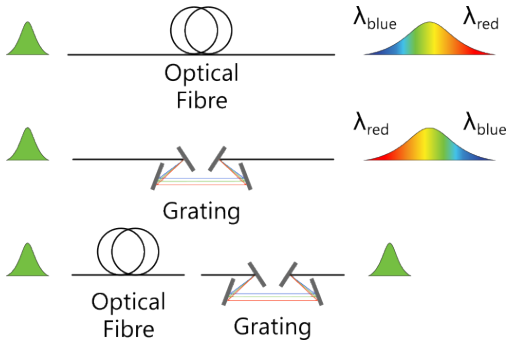

Figure 8: Dispersion compensation with gratings. Top: optical fibre with negative group-velocity dispersion (GVD); Middle: grating compressor with positive GVD; Bottom: combined fibre and grating yielding dispersion-free pulse.

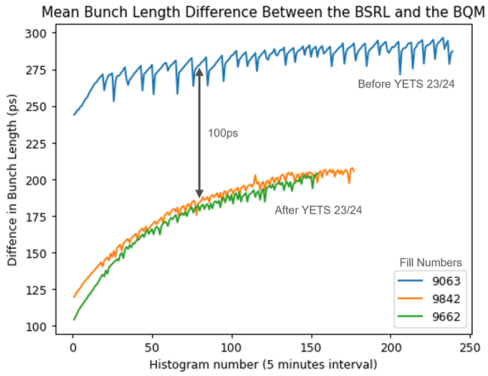

Figure 9: Difference in 4 sigma bunch length measurements from BSRL and BQM fills Before and after the change of timing reference.

this Gaussian is 300 ps, which matches the BST turn clock uncertainty. To reduce this timing error and to profit from the relative closeness of the RF system, the BSRL can take timing signals directly from the RF system. This change was implemented in the 2023-2024 Year End Technical Stop (YETS). The difference can be seen in Fig. 9.

AGEING EFFECTS

Over the course of the LHC Run 3 (5 July 2022 - present), a decrease in count rate was observed as shown in Fig 10. This loss was initially ascribed to deteriorating angular coupling that required manual alignment, which was only possible during shutdowns. A remote angular control device was added in 2024-2025 YETS without resolving the issue, and so the exact location of the losses was sought [12].

In Fig. 7 it is seen that the attenuation of the 100 m long patch fibre from the tunnel to the gallery exhibits an attenuation < 3 dB for each beam, when measured at a wavelength of 1310 nm. However, when injecting 7.5 µW at 532 nm into the same patch fibre, the signal was below the detection limit of the power meter (< 50 nW). In Fig. 10 the count rate partially recovers in the absence of beam, which could be evidence of annealing of the fibre. Taken together the

ISBN: 978-3-95450-262-2 ISSN: 2673-5350 doi: 10.18429/JACoW-IBIC2025-MODC02

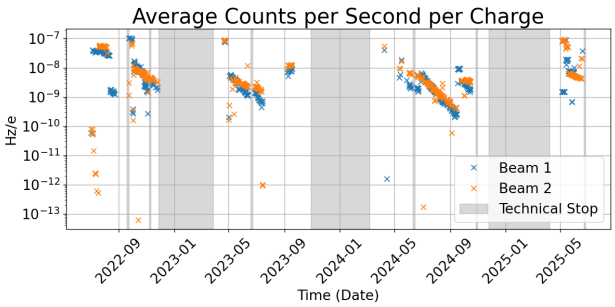

Figure 10: Average photon count rate per proton during stable beams in Run 3 at 6800 GeV. Each point corresponds to one fill per beam, normalized to beam charge.

location of losses, that the effect is less severe with longer wavelengths [13], and annealing point to the loss in counts resulting from Radiation Induced Attenuation (RIA) in the optical fibre.

It is envisioned that both fibres will be replaced in the LHC Long Shutdown 3 (September 2026 - 2028)

Figure 10 shows the photon count rate recorded during stable beams in Run 3, restricted to fills with a duration longer than one hour. To compare fills with different beam intensities, the rate is normalized to the total beam charge. Each data point therefore represents one fill for each beam at 6800 GeV.

At present, the BSRL uses a graded index fibre to limit modal dispersion. This fibre, however, does not have radiation hardening. A possible alternative would be the use of radiation-hard single-mode fibre (SMF), which would eliminate modal dispersion and reduce RIA. However, it is more challenging to achieve and maintain a good light coupling with SMFs, and for this reason, the BSRL moved away from SMFs. Another possibility may be to move the wavelength of the bandpass filter to where the RIA is less severe with longer wavelengths [13].

SUMMARY

This study has shown that the dominant source of counting error in the BSRL is background correction, which is itself sensitive to timing jitter, while the largest timing error stems from chromatic dispersion in the optical fibres transporting synchrotron light. Several mitigation measures have already been implemented, including replacing the Beam Synchronous Timing signal with a direct RF timing reference, greatly reducing jitter, and upgrading to a new time-to-digital converter with ~1 ns dead time and 5 ps precision. Future work will target further improvements through replacement of the current fibres with radiation-hard alternatives and/or adjusting the bandpass filter to longer wavelengths to mitigate radiation-induced attenuation. In parallel, optimising the count rate and integration time with the upgraded system will maximise accuracy, ensuring the BSRL continues to provide reliable measurements for LHC operation and luminosity calibration.

REFERENCES

[1] G. Aad *et al.*, "Beam-induced and cosmic-ray backgrounds observed in the ATLAS detector during the LHC 2012 proton-proton running period", *J. Instrum.*, vol. 11, no. 05, p. P05013, May 2016. doi:10.1088/1748-0221/11/05/P05013

[2] A. Jeff, "A longitudinal density monitor for the LHC", Ph.D. thesis, Phys. Dept., Univ. of Liverpool, U.K., 2012.

[3] M. Gonzalez Berges, A. Jury, and S. Mazzoni, "Beam synchrotron radiation for the longitudinal density monitor (BSRL), error analysis", CERN, Geneva, Switzerland, Rep. CERN-EDMS-3324712, Jul. 2025.

[4] *User Manual: Agilent Acqiris Time-to-Digital Converters, U1092-90017*, Agilent Technologies, Jul. 2012.

[5] PicoQuant GmbH, MultiHarp 150 high-throughput multichannel event timer & TCSPC unit. https://www.picoquant.com/images/uploads/downloads/18277-datasheet_multiharp150.pdf

[6] R. Paschotta, "Photon counting", *RP Photonics Encycl.*, 2007. https://www.rp-photonics.com/photon_counting.html

[7] M. Palm, "BSRL – The LHC longitudinal density monitor", presented at CERN Beam Instrumentation Day, Geneva, Switzerland, 2016, unpublished. https://indico.cern.ch/event/489382/contributions/1166998/attachments/1241315/1825555/BI-Day_BSRL_v3.pdf

[8] K. Wille, "Synchrotron radiation sources", *Rep. Prog. Phys.*, vol. 54, no. 8, p. 1005, Aug. 1991. doi:10.1088/0034-4885/54/8/001

[9] I. H. Malitson, Interspecimen comparison of the refractive index of fused silica. https://refractiveindex.info/?shelf=glass&book=fused_silica&page=Malitson

[10] E. Moreno-García, J. M. de la Rosa-Vázquez, S. Stolik-Isakina, and P. Reyes-López, "Methods for analysis of time-resolved fluorescence data", in *Proc. CONIELECOMP'12*, Cholula, Puebla, Mexico, 2012, pp. 317–322. doi:10.1109/CONIELECOMP.2012.6189931

[11] D. Dominguez, J. J. Gras, J. Lewis, J. J. Savioz, J. Serrano, and F. J. Ballester, "An FPGA-based multiprocessing CPU for beam synchronous timing in CERN's SPS and LHC", in *Proc. ICALEPCS'03*, Gyeongju, Korea, Oct. 2003, pp. 113–115.

[12] D. Butti and A. Schlögelhofer, "BSR meeting – BSRL status", presented at BSR Meeting, Geneva, Switzerland, Jun. 2025, unpublished.

[13] S. Girard *et al.*, "Recent advances on radiation-hardened optical fiber technologies", in *Proc. OFC'20*, San Diego, CA, USA, Mar. 2020, pp. 1–3.