

# ONLINE PERFORMANCE MONITORING AND EARLY FAULT DETECTION FOR THE B-TRAIN REAL-TIME SYSTEMS AT CERN

K. Diamantis\*, M. Bonora-Tam, M. Buzio, V. di Capua, D. Giloteaux, CERN, Geneva, Switzerland

## Abstract

At CERN, B-Train systems provide real-time magnetic field information to the RF and other subsystems with high precision and reliability, with less than 3 beam-hours lost per year per machine. Using the Proton Synchrotron Booster (PSB) as a case study, we describe a software performance-monitoring framework, currently in the prototyping phase, aimed at improving the absolute accuracy and trustworthiness of the measurement to help operators to proactively diagnose a variety of beam instabilities. The framework compares measurements with historical data, model-based expectations, and redundant acquisition chains to assess internal errors due e.g. to integrator drift or field marker degradation, apply the necessary corrections, or flag systems for intervention. We expect it to enhance compliance with evolving operational standards and future large-scale applications requiring automated self-diagnostics and remote maintenance.

## INTRODUCTION

At CERN, six B-train systems provide real-time magnetic field information based on induction coils that generate a voltage  $V_c$  proportional to the derivative of the linked flux, so that by integration we obtain the average bending field  $B(t)$ :

$$B(t) = B_o + \Delta B(t) = B_o - \frac{1}{A_c} \int_{t_0}^t V_c(\tau) d\tau, \quad (1)$$

where  $A_c$  is the effective coil area and  $B_0 = B(t_0)$  is the magnetic field at the beginning of the integration process, signaled by field markers [1]. These are nuclear magnetic (NMR) or ferrimagnetic (FMR) resonance sensors that trigger on each cycle at a preset value, usually just before injection, to reset the integration and minimize the accumulation of drift due to parasitic currents or localized thermoelectric effects [2]. Each system is composed of a redundant pair of acquisition chains, the operational (OP) and the spare (SP), designed to provide identical results. Each chain includes signal conditioning and custom digital electronics hosted in an industrial Front-End Computer (FEC), and takes inputs from different arrangements of sensors, according to the characteristics of the reference dipoles available in each machine.

The different configurations can be described in a unified way via the concept of *logical magnet gap*, i.e. an abstract source of magnetic field driven by a given current. In the simplest cases, the Proton Synchrotron (PS) and Extra Low ENergy Antiproton (ELENA) rings, this corresponds to the physical gap of a reference magnet hosting sensors for both the OP and SP chains. In the Super Proton Synchrotron

(SPS), two dipoles of different types are connected in series and instrumented as a single logical unit. In the Low Energy Ion Ring (LEIR), OP and SP sensors are installed in the fringe field region of two distinct, but nominally identical, bending dipoles in the ring. The Proton Synchrotron Booster (PSB), consisting of four superimposed rings, is the only one where the B-train has two logical gaps, corresponding to the symmetrical pairs of outer (Ring 1 and 4) and inner (Ring 2 and 3) physical gaps, vertically stacked within the same iron yoke. The two outer and the two inner rings are powered in series but independently, to generate identical fields despite different saturation characteristics.

## BTRAIN MONITORING FRAMEWORK

The monitoring framework is organized into six hierarchical levels, shown in Fig. 1. Levels 0 to 4 are per-cycle diagnostic signals computed in real-time, generating warnings when issues are detected. Over 120 diagnostics are defined, many of which are iterated over multiple signal paths per chain. The first three levels define quantities related to any OP or SP chain: low-level voltage signal statistics (Level 0), measured field quality indicators (Level 1) and internal consistency metrics based on the comparison of the measured field with the nominal cycle ( $B_{sim}$ ) and a physical model based on the excitation current (Level 2). Level 3 handles redundancy by evaluating the consistency of OP and SP chains, and Level 4 evaluates consistency across logical magnet gaps. In Level 5 a subset of the computed metrics is stored according to a tiered retention policy for off-line evaluation of historical trends, correlations and discontinuities, either automatically or on-demand, with optional cycle type filtering.

## TECHNICAL IMPLEMENTATION

Implementation of the monitoring framework relies on a multi-tier software architecture designed for real-time data acquisition, processing, and visualization.

### FESA

Front End Software Architecture (FESA) C++ classes run on FECs and provide the means for gathering and filtering the data from the low level hardware connected to magnetic sensors. FESA ensures hardware integrity and provides Level 0 diagnostics. FESA classes are synchronized with the timing system, which distributes trigger pulses across the accelerator complex [3,4]. In our implementation, data are gathered for every triggered event and then distributed through CERN's control middleware (CMW).

\* konstantinos.diamantis@cern.ch

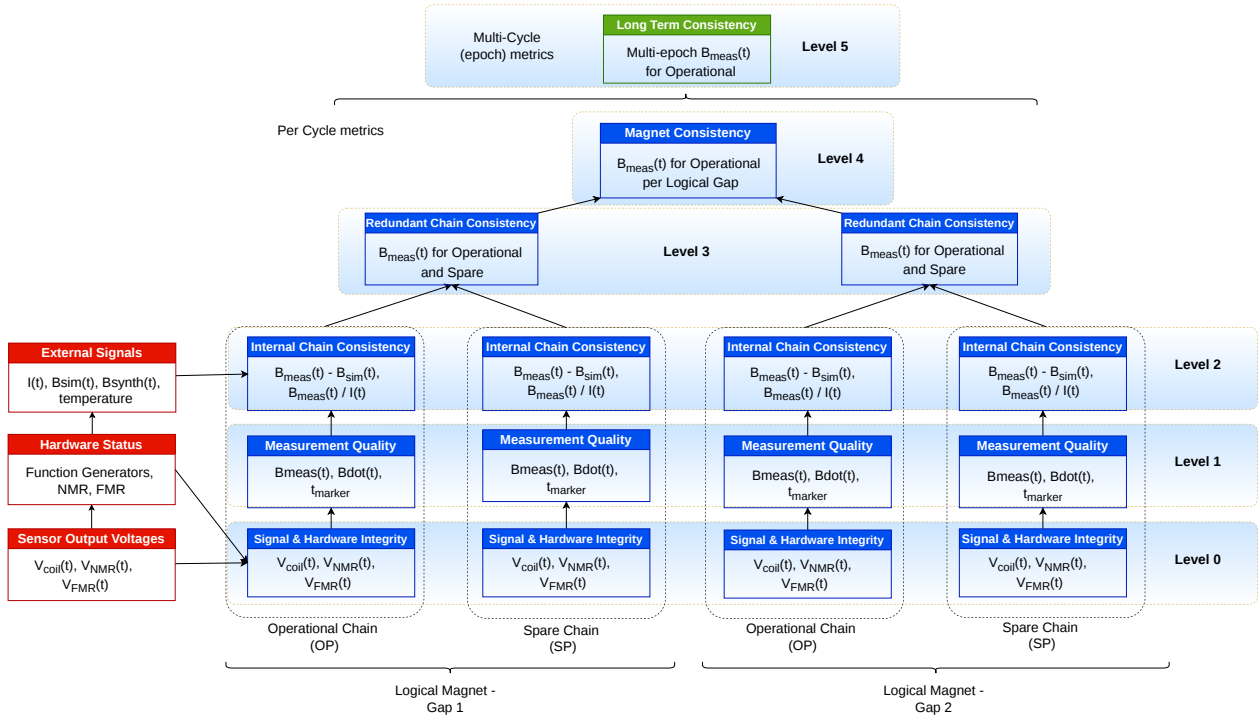


Figure 1: Schematic representation of the different levels of diagnostics and their relationship.

## UCAP

The Unified Controls Acquisition and Processing (UCAP) framework [5] is used in the CERN Accelerator Control System to streamline data acquisition, processing and publishing. A UCAP device can take in real-time data published by different FECs i.e. by different acquisition chains, which is crucial for Level 3 and 4 diagnostics. The framework provides functions for grouping the synchronized data based on timing events, data processing and publishing the results to different clients.

## WRAP & NXCALs

The CERN Web Rapid Application Platform (WRAP) is utilized to construct dedicated graphical user interfaces (GUIs) for real-time visualization [6]. Acting as the consumer of the synchronized data published by UCAP on dedicated nodes, WRAP organizes all critical monitoring metrics from B-Train acquisition chains into centralized dashboards. WRAP subscribes directly to live data feeds from CMW, enabling users to monitor the health of the system in real-time [7]. The data streamed to WRAP can be stored in the Accelerator Logging Service (NXCALs) [8]. This enables Level 5 diagnostics by executing scheduled Python-based analysis scripts that feed status updates to the same WRAP instance. This visualization strategy will help identify performance anomalies and track long-term drifts.

## RESULTS AND DISCUSSION

We demonstrate the prototype implementation of the B-Train monitoring framework focusing on the Ring 2 and 3 of the PSB. A WRAP Dashboard example is shown in Fig. 2

while a subset of diagnostic metrics is given in Table 1. For every magnetic cycle, key parameters are aggregated over a rolling window, 50 cycles wide in this example. Results are displayed in the Dashboard and their average and standard deviation are given in the last column of Table 1. Metrics include:

- The absolute maximum and minimum field  $B$ , to be compared with the extrema expected on the basis of the well-known magnet transfer function. The observed minimum is slightly negative, due to a small power converter undershoot.
- The marker trigger state, i.e. a boolean value that indicates whether the FMR marker's output voltage  $V_{FMR}$  crosses the negative detection threshold  $\bar{V}_{FMR} = -0.15$  V. Sporadic misses (such as in the example shown, where 90% of the cycles in the analysis window registered a trigger) raise a visual flag in the Dashboard but do not impact functionality, while repeated contiguous misses raise a high-priority alarm.
- The minimum marker output  $\text{Min}(V_{FMR})$  and the corresponding detection margin, defined as  $\text{Min}(V_{FMR}) / \bar{V}_{FMR} - 1$ . In time,  $V_{FMR}$  may attenuate due to sensor aging, temperature drift or shifts in the baseline noise floor, and the resulting long-term margin erosion can be tracked to raise a warning. Here the average peak output is very strong, almost double the threshold, which suggests that the cause of missed triggers is not related to signal levels.

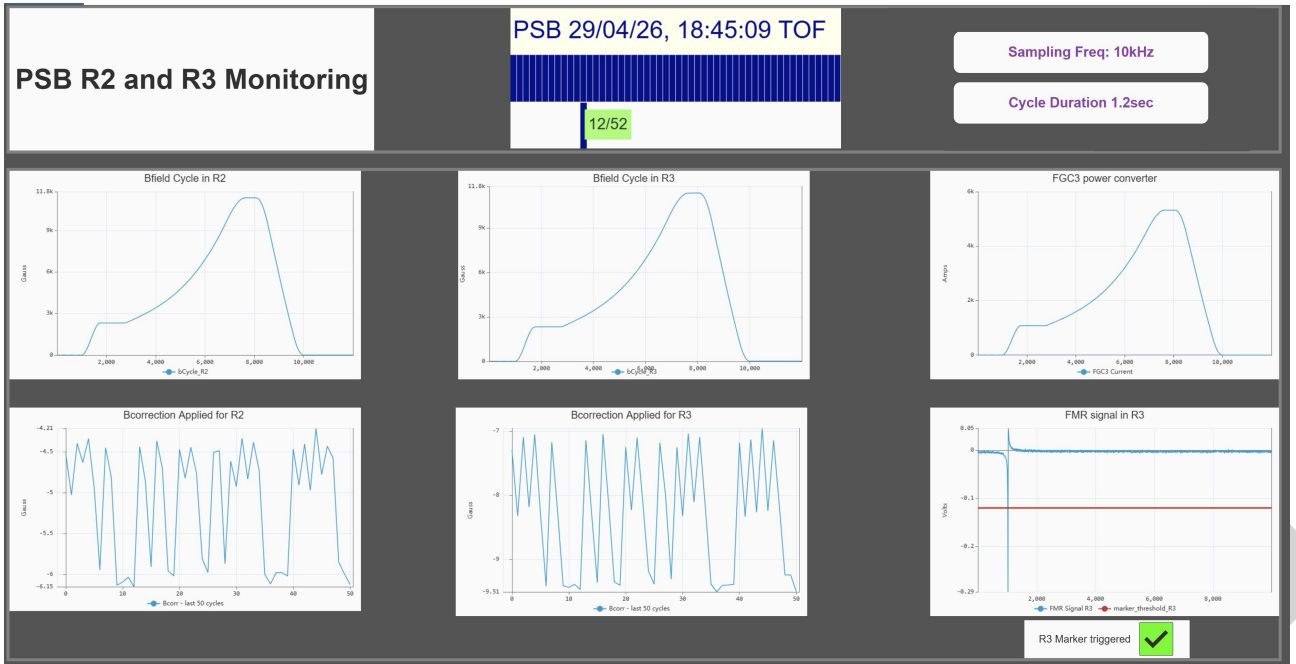


Figure 2: WRAP Dashboard for PSB Rings 2 and 3. Displays the last played cycle metrics, including timing, magnetic fields, power converter current, field corrections applied during last 50 cycles and Ring 3 FMR signal. Additional real-time analytics are accessible via supplementary Dashboards of the same WRAP instance.

Table 1: Examples of Diagnostic Metrics for the PSB R2 Ring

Level	Metric (Unit)	Expected	Observed <sup>*</sup>
0	Min ( $V_{\text{FMR}}$ ) (V)	$< -0.15$	$-0.29 \pm 0.02$
1	Max( $B$ ) (T)	1.29	$0.96 \pm 0.11$ **
1	Min( $B$ ) (mT)	0	$-0.55 \pm 0.14$
1	Marker trigger state	100% triggered	90% triggered
1	$t_{0,\text{R2}}$ (ms)	135	$135.0 \pm 0.1$ **
1	$B(t_0) - B_0$ (mT)	0	$-0.51 \pm 0.07$
2	Marker detection margin	$> 10\%$	$93.3\% \pm 13.3\%$
2	$\dot{B}_{\text{drift}}$ (mT s <sup>-1</sup> )	$ \cdot  < 0.1$	$-0.455 \pm 0.04$
3	$t_{0,\text{R2}} - t_{0,\text{R3}}$ (ms)	$ \cdot  < 1$	$4.0 \pm 0.1$ **

<sup>\*</sup> Average  $\pm$  standard deviation of Level 5 diagnostic of the given metric, calculated over a typical 50-cycle window

<sup>\*\*</sup> Considering only TOF cycles (interleaved ZERO cycles excluded)

- The cycle time at which the R2 marker trigger is received,  $t_{0,\text{R2}}$ . This is stable as expected, within the sampling interval tolerance.
- Field correction, i.e. the difference between the running integral at the time of a marker trigger  $B(t_0)$  and the marker preset value.
- The average field drift rate  $\dot{B}_{\text{drift}}$ , defined as the field correction divided by the time interval between the current and last marker trigger. The high observed values show the need for the physical recalibration of the sensor.
- The delay between marker triggers in the R2 and R3 rings, which should arrive at the same time. The observed 4 ms delay is stable within the sampling interval tolerance, which suggest to look into a systematic calibration difference of the two markers.

## SUMMARY AND OUTLOOK

We developed a scalable, multi-level diagnostic and monitoring software framework for CERN B-Train systems. It represents a radical shift from reactive interventions to a preemptive approach toward the automatic detection of subtle, progressive system degradation as well as faults requiring urgent attention. By leveraging software modularity, we successfully integrated real-time acquisition with diagnostic algorithms that monitor hardware integrity, signal quality, and measurement consistency. After this initial operational validation, the present prototype will be extended to all machines to improve the reliability of these crucial measurement systems in the future.

## REFERENCES

- [1] M. C. L. Buzio *et al.*, “Development of upgraded magnetic instrumentation for CERN’s real-time reference field measure-

- ment systems”, in *Proc. IPAC'10*, Kyoto, Japan, May 2010, paper MOPEB016, pp. 310–312. <http://accelconf.web.cern.ch/IPAC10/papers/MOPEB016.pdf>
- [2] C. Grech, M. Buzio, and N. Sammut, “A magnetic measurement model for real-time control of synchrotrons”, *IEEE Trans. Instrum. Meas.*, vol. 69, no. 2, pp. 393–400, 2020. [doi:10.1109/TIM.2019.2904073](https://doi.org/10.1109/TIM.2019.2904073)
- [3] A. Topaloudis and C. Rachex, “Visualisation of real-time front-end software architecture (FESA) developments at CERN”, in *Proc. ICALEPCS'17*, Barcelona, Spain, Oct. 2017, pp. 1853–1856. [doi:10.18429/JACoW-ICALEPCS2017-THPHA180](https://doi.org/10.18429/JACoW-ICALEPCS2017-THPHA180)
- [4] A. Schwinn *et al.*, “FESA3: the new front-end software framework at CERN and the FAIR facility”, in *Proc. PCaPAC'10*, Saskatoon, Canada, Oct. 2010, paper WECOAA03, pp. 22–26. <https://jacow.org/pcapac2010/papers/WECOAA03.pdf>
- [5] L. Cseppentő and M. Büttner, “UCAP: a framework for accelerator controls data processing at CERN”, in *Proc. ICALEPCS'21*, Shanghai, China, pp. 230–235, Mar. 2022. [doi:10.18429/JACoW-ICALEPCS2021-MOPV039](https://doi.org/10.18429/JACoW-ICALEPCS2021-MOPV039)
- [6] E. Galatas, A. Asko, E. Matli, and C. Roderick, “WRAP – a web-based rapid application development framework for CERN’s controls infrastructure”, in *Proc. ICALEPCS'21*, Shanghai, China, pp. 894–898, Feb. 2022. [doi:10.18429/JACoW-ICALEPCS2021-THPV013](https://doi.org/10.18429/JACoW-ICALEPCS2021-THPV013)
- [7] E. Galatas, K. Tsiaras, S. Deghaye, and C. Roderick, “WRAP: integrating an event processing framework”, in *Proc. ICALEPCS'25*, Chicago, IL, USA, pp. 1763–1769, Dec. 2025. [doi:10.18429/JACoW-ICALEPCS2025-THPD077](https://doi.org/10.18429/JACoW-ICALEPCS2025-THPD077)
- [8] J. P. Wozniak and C. Roderick, “NXCALs - architecture and challenges of the Next CERN Accelerator Logging Service”, in *Proc. ICALEPCS'19*, New York, NY, USA, pp. 1465–1469, Aug. 2020. [doi:10.18429/JACoW-ICALEPCS2019-WEPHA163](https://doi.org/10.18429/JACoW-ICALEPCS2019-WEPHA163)