Nuclear Instruments and Methods in Physics Research A 1086 (2026) 171244

Contents lists available at ScienceDirect

Nuclear Inst. and Methods in Physics Research, A

journal homepage: www.elsevier.com/locate/nima

Full Length Article

Real-time monitoring of longitudinal beam quality across the CERN accelerator complex

Amaury Beeckman [a,b,*], Alexander Huschauer [a], Alexandre Lasheen [a], Cédric Hernalsteens [a,b], Nicolas Pauly [b]

[a] CERN, European Organization for Nuclear Research, Esplanade des Particules 1, 1211 Meyrin, Switzerland
[b] Service de Métrologie Nucléaire (CP165/84), Université libre de Bruxelles, Avenue Franklin Roosevelt 50, 1050 Brussels, Belgium

ARTICLE INFO

Keywords:
Automation
Beam performance
FESA
UCAP
Wall current monitor
Longitudinal beam observation
Tomography

ABSTRACT

Acquiring and processing longitudinal beam profiles in the CERN injector complex is essential for maintaining high beam quality for both the Large Hadron Collider and fixed-target experiments. This paper describes a new and common framework across the injector chain to continuously monitor longitudinal profiles, based on digitizers connected to wall current monitors. A new software class has been developed using the Front-End Software Architecture framework to improve real-time monitoring of longitudinal beam quality along the whole beam production cycle. The new implementation supports efficient data acquisition with highly flexible configuration options, including per-profile and per-cycle settings management. This acquisition class reduces data transfer time from the hardware and enables simultaneous beam monitoring and optimization for multiple users across the CERN accelerator complex. Beyond the raw data acquisition layer, real-time processing analyses have been developed on the Unified Controls Acquisition and Processing platform to perform bunch length, bunch-by-bunch intensity, and bunch spacing measurements, and tomographic reconstruction of the longitudinal phase space to provide key longitudinal parameters such as momentum spread and emittance. This new monitoring framework paves the way for preparing and executing the filling of the LHC in an automated manner and for optimizing fixed-target cycles for physics users.

1. Introduction

The CERN accelerator complex comprises a variety of linear accelerators and synchrotrons that are interconnected and used to accelerate various species of particle beams. The versatility of the complex enables the beams to be accelerated to different energies, ranging from keV to TeV scales. Apart from the Large Hadron Collider (LHC), the synchrotrons are connected to their dedicated experimental fixed-target facilities, such as the Isotope mass Separator On-Line (ISOLDE) facility and the EAST and NORTH experimental areas [1–3]. The different user facilities and the different synchrotrons have dedicated beam parameter requirements. To deliver beams to all users with a maximum throughput, the complex operates on a supercycle basis [4], a repeating sequence of different beam cycles (comprising injection, acceleration, and extraction) in each synchrotron. The machines must alternate between distinct settings to deliver beams of varying energies, intensities, and time structures. Ensuring beam delivery according to the specifications has always been of the highest importance to reach the defined performance goals and ensure a successful and rich physics program.

Monitoring beam performance has increased in importance during the last decade to meet the specified beam parameter requirements. The needs of the user community and the upgrade programs to extend the performance reach (e.g. doubling the intensity of the beam in view of the LHC upgrade [5]) of the accelerator complex have pushed the necessity for a real-time monitoring system across machines. This real-time monitoring system should measure longitudinal and transverse beam quality during operation to ensure high beam quality for both fixed-target experiments and the LHC at CERN.

This paper will focus only on the longitudinal plane. In a synchrotron, the radio frequency (RF) cavities determine the longitudinal structure of the particle beam according to the frequency of the electromagnetic field in the cavities. These cavities typically operate in the megahertz range in the CERN accelerator complex. Along the CERN accelerator chain, the beam goes through many RF manipulations [6] to change the longitudinal beam characteristics. These RF manipulations, such as bunch splitting, batch compression, and bunch rotation, are necessary to produce the required beam parameters for the LHC and the various experimental facilities and to transfer the beams from one

* Corresponding author at: CERN, European Organization for Nuclear Research, Esplanade des Particules 1, 1211 Meyrin, Switzerland.
E-mail address: amaury.beeckman@cern.ch (A. Beeckman).

https://doi.org/10.1016/j.nima.2025.171244
Received 31 October 2025; Received in revised form 10 December 2025; Accepted 19 December 2025
Available online 22 December 2025
0168-9002/© 2026 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

synchrotron to the next. Longitudinal beam observation relies on the fact that a beam passing inside the vacuum chamber induces an image current that is related to the number of particles in the beam and its time structure. Beam instrumentation devices such as resistive wall current monitors (WCMs) [7] can be used to measure these image currents. By connecting fast digitizers to the WCM, the analog signal is converted into a quantized digital signal, enabling measurement of the longitudinal characteristics.

Over the past 20 years, the standard solution to acquire longitudinal beam profiles in the Proton Synchrotron Booster (PSB) and the Proton Synchrotron (PS) relied on the Open Analogue Signal Information System (OASIS) acquisition system [8], which gathered a sequence of acquisition profiles from a WCM over a specific time window. This raw data is then processed by an operational software application [9,10] to derive the bunch length and longitudinal emittance on demand. OASIS is a low-level software layer in the Front-End Software Architecture (FESA) [11] framework. It provides a generic interface for many digitizer devices and gives an output resembling an oscilloscope. Even if the signals originate from physically distant sources, this virtual oscilloscope allows their data to be displayed on a common graphical user interface in the control room. However, OASIS was not designed for permanent acquisition, which prevents continuous monitoring and control of the synchrotron based on raw data acquisition. Furthermore, it lacks cycle-by-cycle settings management capabilities, which prevents effective monitoring and optimization of beam quality across all beam types simultaneously. In the Super Proton Synchrotron (SPS), stringent interlocking requirements led to the development of a dedicated software acquisition system [12], which triggers a beam dump if the longitudinal beam quality falls outside predefined limits for injection into the LHC. This protection system runs continuously, capturing a limited number of beam profiles at key points in the cycle (injection and extraction) to verify compliance. Due to the criticality of this system, it does not provide any flexibility in its acquisition, preventing it from being used for other beam types and to see longitudinal beam dynamics.

In this paper, two significant advancements related to the monitoring of longitudinal beam quality are described. First, we detail the development of a new software layer in the FESA framework to acquire raw data throughout the accelerator chain. The new acquisition layer ensures continuous and reliable data streams, optimized to acquire up to 100 times more data than previous systems. To meet real-time constraints, hardware configuration, pre-processing, and post-processing is completed within a few hundred milliseconds before each beam injection. It supports per-cycle configuration with multiple time windows and maximum flexibility in settings management, enabling operation across diverse accelerators and beam types. It is designed generically so that it can be deployed in broader contexts, including other fast digitizer applications at CERN and in other facilities [13] using the same control system. Secondly, we present a new infrastructure that computes and monitors key performance indicators of longitudinal beam quality in real-time, such as bunch length, bunch-by-bunch intensity variations, bunch spacing, and tomographic reconstruction of the longitudinal phase space. All these metrics, computed only on demand previously, are now available in real-time through the CERN controls infrastructure and can be visualized in the accelerator control room. By standardizing longitudinal data acquisition, signal processing, and the computation of beam observables, the system provides a unified and reliable framework for longitudinal beam diagnostics. Together, these developments establish the foundation for systematic beam quality tracking and autonomous beam optimization.

The structure of the paper is as follows: Section 2 describes the new acquisition system design, detailing the class processes, user parameter configurations, and the acquisition data fields. The potential of this new acquisition layer is shown in Section 3 by showcasing its ability to acquire longitudinal profiles in different acquisition scenarios. Section 4 introduces the real-time monitoring system that automatically computes the main longitudinal beam quality metrics from the acquisition data. Finally, Section 5 highlights several use-cases made possible by this new global framework for longitudinal beam observation.

2. Raw data acquisition system in FESA

Real-time systems at CERN are controlled by Front-End-Computers (FECs) [4]. These FECs are running FESA classes that enable low-level control of the hardware parameters and expose the raw data acquisitions and configuration properties to higher software layers (e.g. an operational application). The LHC Software Architecture (LSA) [14] database and settings management system handles the hardware parameters and stores the machine setting configurations necessary to produce the different beams. Since physics users have different beam requirements, each beam production cycle has its own machine parameter configuration stored in LSA (evolution of the main magnetic field over time, settings of auxiliary magnets, RF elements, beam instrumentation parameters, etc.). Different cycles are being played back-to-back in the supercycle, with a minimum cycle length of 1.2 s and a maximum length that depends on the accelerator. The FEC, through FESA [11], is responsible for linking the hardware systems with the timing system to prepare the execution of the next cycle by driving the corresponding settings from LSA to the hardware. This cycle-by-cycle settings management is called multiplexing or pulse-by-pulse modulation (PPM).

The FESA framework provides a common integrated coding environment for the development and deployment of real-time control software. It is used to control most of the beam instrumentation devices at CERN, including beam loss monitors [15], beam position monitors [16], and beam current transformers [17]. Each type of measurement device requires a specific FESA class. FESA classes play a role equivalent to EPICS input/output controllers; they expose an interface that higher-level control systems can use to read or modify hardware parameters and to retrieve measurement data from the instrument.

The acquisition system setup and logic to acquire longitudinal beam profile information are shown in Fig. 1. The timing system sends a signal to all FECs to notify them that a new cycle is starting. All FECs, via their FESA classes, can react to this event by loading the specific configuration settings required for that cycle to the hardware devices. In the current system design, a WCM is installed in the accelerator tunnel and measures a continuous analog signal, which is transferred via cables to digitizer cards installed in the FECs outside of the tunnel. The digitizer is an analog-to-digital converter that begins recording when it receives a trigger signal. From that point, it samples the incoming analog signal at regular time intervals, converting each measured point into a digital value representing the signal amplitude. The digitizer card is connected to the FEC via a PCI Express bus, and the data is transferred to the FEC memory via a generic driver interface described below. To start the acquisition, the triggers are generated by a synchronized generator, which is in phase with the revolution frequency of the beam. This synchronization is crucial for longitudinal beam observation, as the signal can then be aligned turn by turn. The digitizer and trigger generator configuration settings for the acquisition are stored in the LSA database per beam cycle. These settings are synchronized with the FESA settings interface to ensure a consistent hardware setup. Once the acquisition process is complete, the data is made available over the network to downstream clients, who can further process and display it.

The following sections detail the implementation of a new raw data acquisition FESA class designed to overcome the OASIS limitations. The acquisition layer is specifically designed for longitudinal beam observation using fast digitizers, enabling a cycle-by-cycle configuration with a permanent data acquisition stream.

2.1. Internal design of the FESA class

The presented FESA class has been primarily developed to acquire data from Acqiris digitizers across most CERN synchrotrons, including the PSB, the PS, the SPS, and the LHC. Different digitizer models with varying acquisition sampling rates (e.g. SA248P and SA217P [18,19])

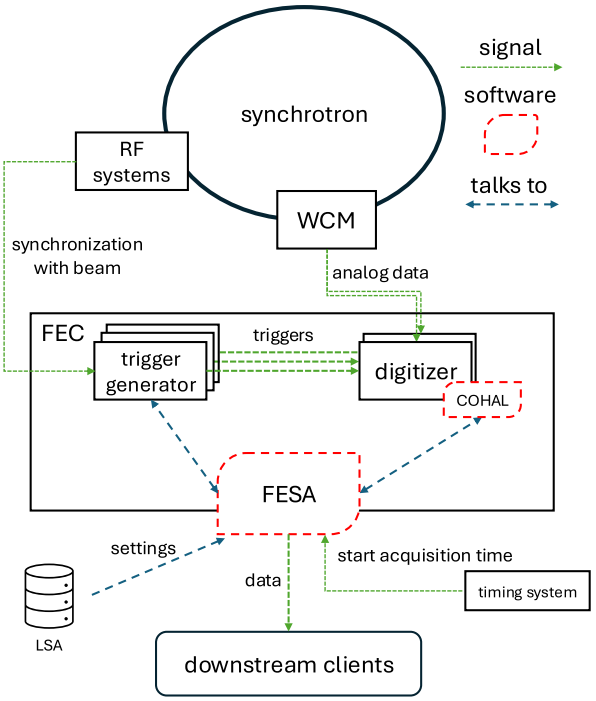

Fig. 1. Schematic view of the acquisition system setup for longitudinal beam profiles. The timing system synchronizes the different hardware instruments via the FESA framework with the start-of-acquisition timing event. The digitizer receives the analog signal from the WCM and starts acquiring profiles with trigger signals from a trigger card. The trigger card is synchronized with the revolution period of the beam, i.e. with the RF systems. COHAL is the abstraction library that standardizes the communication to the hardware (see Section 2.1). The settings are stored in the LSA database, from where they are propagated to the hardware. Once the acquisition is complete, the data is sent over the network to the downstream clients.

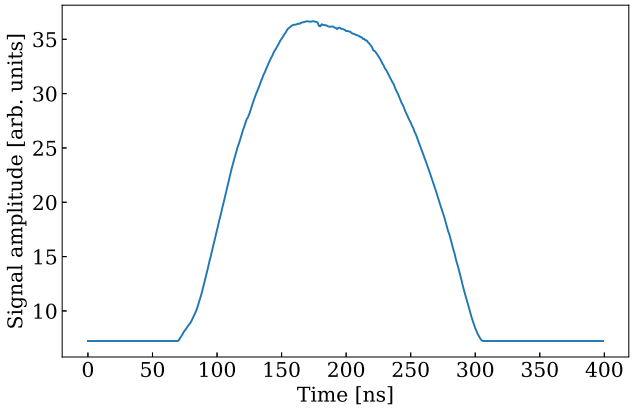

Fig. 2. One bunch profile measured by the digitizer at the extraction energy in the PSB. The digitizer acquires raw data by sampling at 1 GS/s. A total of 400 data points were collected to accurately measure the bunch length (~220 ns corresponds to 220 points).

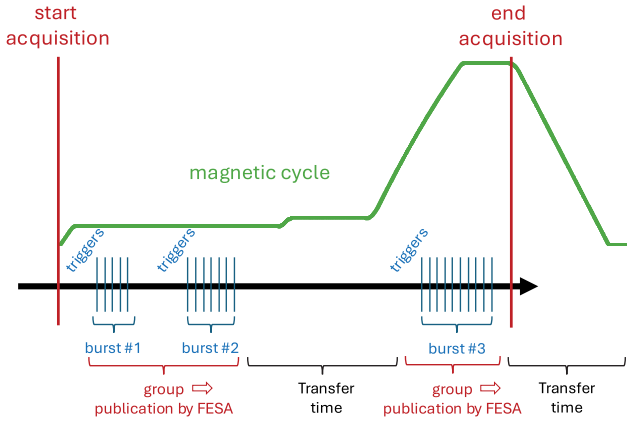

Fig. 3. The figure shows a schematic view of the acquisition process along a typical cycle. The digitizer acquires profiles between the start- and the end-of-acquisition timing. A trigger card generates the incoming triggers in a burst. A burst is an ensemble of triggers generated with a specific delay. The groups correspond to several bursts published together by the FESA class. The first group is acquired at injection energy, and the second is acquired at the end of acceleration and on the extraction flat top. The end-of-acquisition timing is linked to the end of the flat top of the magnetic cycle.

are used at CERN, with their application deployment being based on the beam properties during acquisition. Fig. 2 illustrates the data acquisition of a single bunch at extraction energy in the PSB. The bunch length is approximately 220 ns, and the digitizer collects data at a sampling rate of 1 GS/s. This digitizer configuration enables the acquisition of 220 data points to analyze the bunch properly, which is ideal to quantify its longitudinal structure. However, the same configuration would not be adapted for the LHC, where the bunch length is approximately 1.2 ns, meaning that only one point would be gathered. The LHC uses digitizers operating at sampling rates of up to 40 GS/s. To abstract the specificity of the digitizers, the FESA class relies on the CERN-internal Controls Hardware Abstraction Layer (COHAL) [20] software library. This library standardizes the communication between FESA and the hardware instrument.

The digitizer acquires profiles between the start- and the end-of-acquisition event. These events can be controlled by timing events or by on-demand commands. For fast-cycling machines, acquisition begins with the start cycle timing event and ends with the extraction timing event, as shown in Fig. 3. The starting event will configure the hardware settings for the incoming acquisition, and the end event will stop the acquisition if the digitizer is still waiting to acquire beam profiles. In the LHC, where stable beam conditions can last for hours, start and stop commands can also be issued on demand for greater flexibility.

The digitizer acquires one beam profile per received trigger, typically corresponding to a full turn in which all bunches are measured. Triggers are usually generated in *bursts*, each consisting of N triggers in succession with a defined delay between them as illustrated in Fig. 3. As the trigger generator is synchronized with the beam revolution period, the evolution of the bunch profile can be reconstructed by concatenating the individual profiles, as illustrated in Fig. 4. The longitudinal synchrotron motion of the beam can therefore be observed over the N acquired profiles.

At the start-of-acquisition event, the FESA class calculates the maximum number of required samples from user settings. This maximum number depends on the number of samples, the acquisition sampling rate, and the trigger delay of the most demanding configuration for the entire acquisition. If the sampling rate selected by the user is smaller than the sampling rate of the digitizer hardware, an averaging is applied instead of a decimation to reduce the acquisition noise from the signal. The digitizer is then configured with these settings, armed, and waits for all profiles to be acquired. Once the acquisition is complete, the FESA class directly retrieves the raw data from the digitizer and performs post-processing before publishing the results.

The FESA class publishes the data per *group* as shown in Fig. 3, where each group can include one or more bursts. The data for the

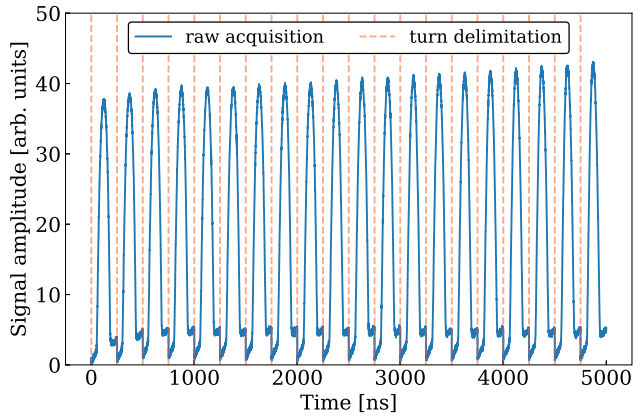

(a) Data acquisition signal recorded by the digitizer during several consecutive turns, where the digitizer acquired 500 points per turn. The dashed red lines delimit individual turns.

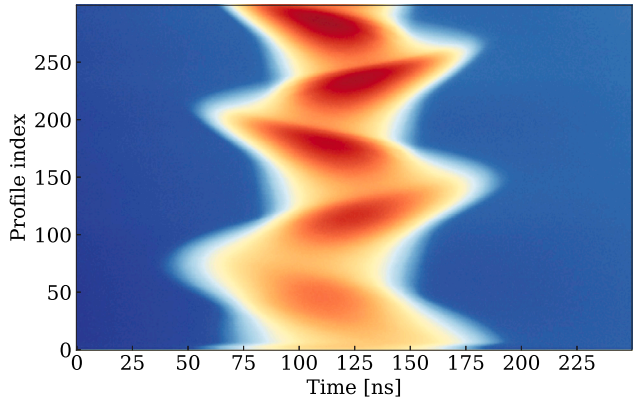

(b) The stacking of all the turn profiles constitutes a waterfall plot, where the synchrotron oscillations can be observed. The color scale represents the signal amplitude, with blue as the background and red as the maximum.

Fig. 4. Data acquired by one digitizer in the PS at injection.

first group is acquired and published, with the subsequent groups following the same principle. This is particularly useful for long cycles, as metrics can be computed as soon as each group is acquired rather than waiting for the full acquisition. This process continues until all user-programmed settings have been processed or the end-of-acquisition event occurs. The end-of-acquisition event stops the acquisition process, preventing the acquisition of unwanted profiles from the next cycle.

As mentioned, the COHAL library standardizes communication with hardware instruments. For most devices acquiring longitudinal profiles, digitizers from the manufacturer allow for bulk reading (read the whole memory buffer in one call) to improve communication efficiency, removing the overhead of reading profile by profile from memory. This crucial functionality was initially missing in COHAL and was therefore added in the framework of the work presented in this paper. This optimization decreased the read-out time by at least a factor of four (see Fig. 5), enabling a transition from profile-by-profile reading to a full memory read in a single function call.

2.2. Acquisition modes

The FESA class was developed to handle various use cases, covering the needs for monitoring operational beams as well as more specific acquisitions for machine development studies. It supports applications for both fast-cycling machines, with repetition rates of a few seconds, and

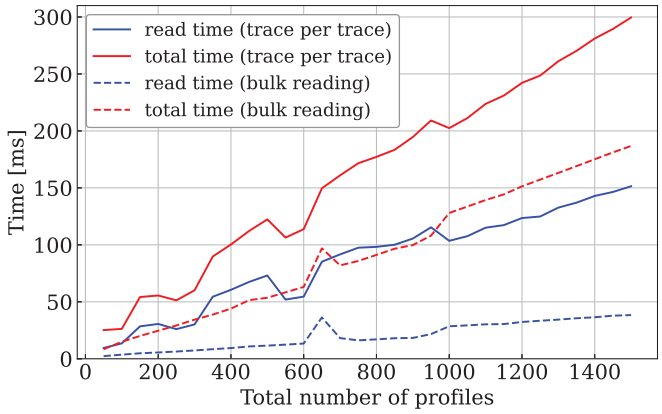

Fig. 5. Performance gain in time for the acquisition process due to code optimization in the COHAL library that incorporates bulk reading. The total time, including publication, pre- and post-processing times, is shown in red, and the reading time is shown in blue. The reading time was reduced on average by a factor of 4.

very long cycles of several hours, as in the LHC. As the digitizer usage varies, three main acquisition modes were designed to accommodate these differences: burst, trace, and continuous acquisition modes.

2.2.1. Burst acquisition mode

The burst acquisition mode is designed for operational use cases and quick setup scenarios in fast-cycling machines. In this mode, all profiles within a burst share identical settings. This common configuration simplifies the setup of multiple profile acquisitions. Once configured, the FESA class publishes all groups defined by the user and then waits for the next start-of-acquisition event.

2.2.2. Trace acquisition mode

The trace acquisition mode allows the specification of individual settings for each beam profile, offering greater flexibility. This acquisition mode is beneficial for acquisitions requiring adjustments based on profile-by-profile variations of the revolution frequency, such as during particle acceleration and orbit bumps. The profiles can be aligned by defining different delay configurations for each profile, as shown in Fig. 6, facilitating the subsequent analysis. With this configuration flexibility, the trace acquisition mode is also ideal for applications where the user requires full control of the parameters for each acquisition profile. Similar to the burst acquisition mode, the FESA class publishes all user-defined groups once the acquisition is complete and waits for the next start-of-acquisition event.

2.2.3. Continuous acquisition mode

As for the trace mode, the continuous acquisition mode defines settings per profile. However, the digitizer is configured for a single burst of triggers that repeats continuously until the end-of-acquisition event arrives. After acquiring all the profiles from one burst, the FESA class publishes the data immediately and automatically rearms the digitizer with the same configuration. This mode is primarily used to monitor the profile evolution over extended periods and is well-suited for slow-cycling machines like the SPS and LHC.

2.3. FESA class settings management

The FESA class supports multiplexed configuration, allowing settings to differ between cycles. Its configuration interfaces enable users to configure device and acquisition properties.

Device settings refer to the hardware configuration of the digitizer. Changes to these settings may require the digitizer to recalibrate.

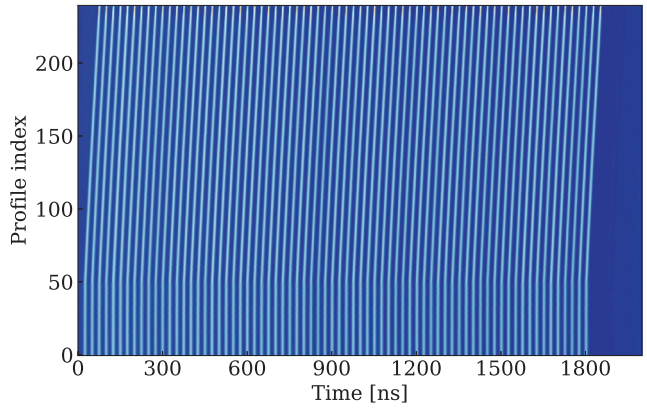

(a) Profiles acquired without considering the variation of the revolution frequency (burst mode).

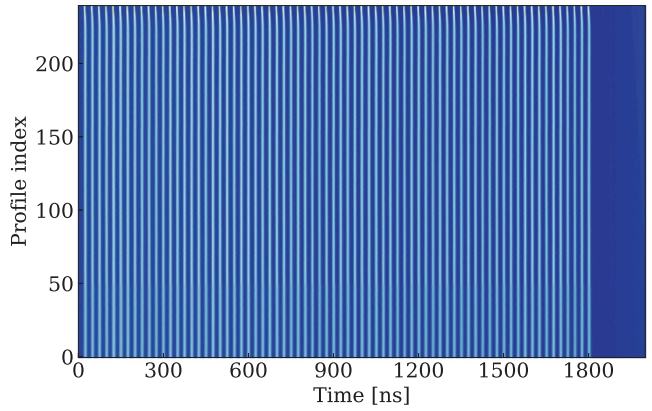

(b) Profiles acquired with a correction of the delay per beam profile to account for the revolution frequency change induced by the extraction bump (trace mode).

Fig. 6. Data acquisition in the PS on a cycle with beam parameters corresponding to the LHC beam requirements, during the rise of the extraction bump.

Table 1
Available parameters in the acquisition settings of the FESA class depending on the acquisition mode (burst, trace, and continuous modes).

	Burst	Trace	Continuous
Group ID	✓	✓	✗
Burst ID	✓	✓	✗
Number of samples	✓	✓	✓
Number of profiles	✓	✗	✗
Sample interval (ns)	✓	✓	✓
Trigger delay (ns)	✓	✓	✓

The new generation of Acqiris digitizers used for longitudinal beam observation features smart calibration in their drivers, which performs recalibration only when a specific set of settings is applied for the first time. The configurable device settings include the vertical sensitivity range, the digitizer acquisition sampling rate, and the vertical signal offset.

Acquisition settings, on the other hand, are directly related to the configuration of the acquisition process. The available properties in the acquisition settings of the FESA class depend on the acquisition mode, as shown in Table 1. These settings are valid for an entire burst (in burst mode) or individual profiles (for trace and continuous modes). In the FESA class, each acquisition setting is defined as a sequential list

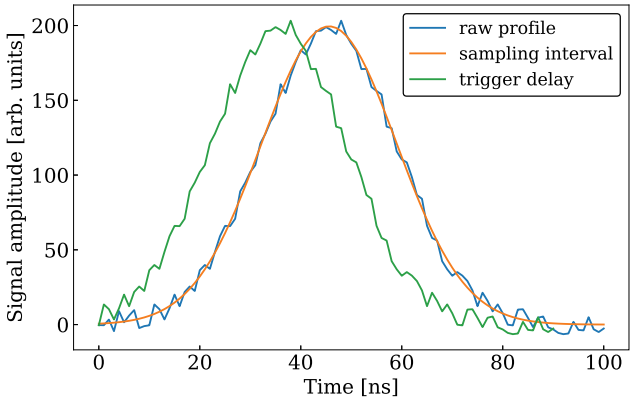

Fig. 7. Impact of different acquisition settings of the FESA class on the measured profile. The blue profile corresponds to the raw acquisition profile, including some noise. The averaging window smooths the raw curve and removes the noise by modifying the acquisition sampling interval, as shown in orange. The trigger delay allows the raw acquisition to be shifted in time, as shown in the green profile, reducing the number of points needed if the bunch is far from the trigger timing.

along the cycle. All parameters at a given index correspond to the same profile or burst.

The user can define groups of profiles published by the FESA class by specifying the group identifiers. Another acquisition setting is the software sampling interval, which defines the software averaging applied during post-processing. The ratio between this acquisition setting and the digitizer sampling interval gives the averaging window, allowing for the removal of the profile noise. The trigger delay specifies a software delay before the first acquisition point. The device will acquire more points starting from the trigger, and the first N points corresponding to this delay will be removed in post-processing. The impact of the different acquisition settings on the profile is shown in Fig. 7.

The real-time requirements of cycle-by-cycle acquisition impose performance constraints on the FESA class, as digitizer readout and raw data post-processing must be completed within a few hundred milliseconds before the start of a new cycle. Reading the digitizer data accounts for roughly one-third of the total processing time required by FESA to publish the data. A higher number of samples and trigger delay settings directly increase the volume of data that must be transferred from the digitizer to the FEC. Among the acquisition settings, the sampling interval has the largest impact on performance, as it generates proportionally more data points and requires post-processing to average them, which constitutes the main bottleneck for time performance. To mitigate this, the software limits the averaging to a maximum of four points, balancing data resolution with processing efficiency. Finally, network bandwidth imposes an additional constraint with a 1 Gb/s link. The amount of data that can be transmitted through the publication is therefore strictly limited.

2.4. Publication of acquired data

For convenience, the class publishes all acquisition settings, including the number of samples and profiles, along with the acquired data and metadata related to the acquisition. These properties are always published per profile, regardless of the digitizer acquisition mode.

The acquisition data is formatted as a 1D array, with all profiles concatenated for efficiency. The user is then responsible for reshaping the data according to the number of samples per profile. This design was chosen to support the flexibility of different acquisition settings per profile, making it the most efficient way to transmit data over the network. This approach avoids the need for padding when different numbers of samples are acquired per profile.

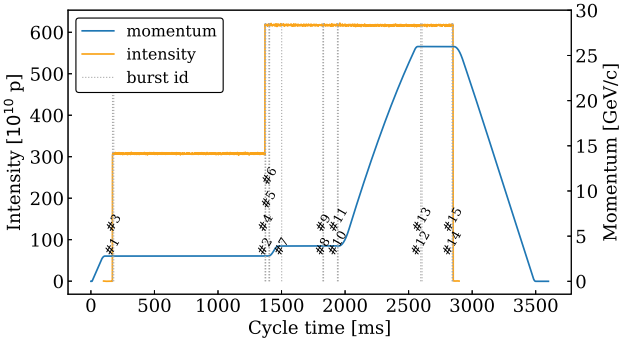

(a) Evolution of the magnetic field and the intensity along the cycle. The two injections are clearly highlighted by the increase in intensity. The vertical gray lines represent the start time of each burst.

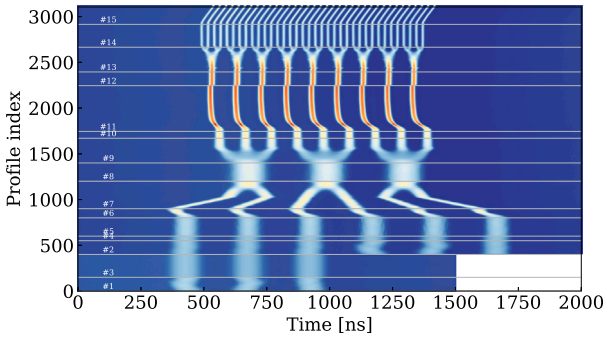

(b) Waterfall representation of the measurement from injection to extraction. The x-axis corresponds to the time in nanoseconds per profile acquired by the instrument, while the y-axis shows the index of the different profiles along the cycle. The color in the figure represents the signal amplitude in arbitrary units. Each burst is separated by a gray line in the plot. The batch compression and the merging can be observed in burst number 7, while the splittings can be seen in the bursts 9 and 13.

Fig. 8. Longitudinal beam profile acquisitions on an LHC BCMS beam produced in the PS. For these acquisitions, the burst timings nearly overlap to ensure that the measurements cover the whole cycle.

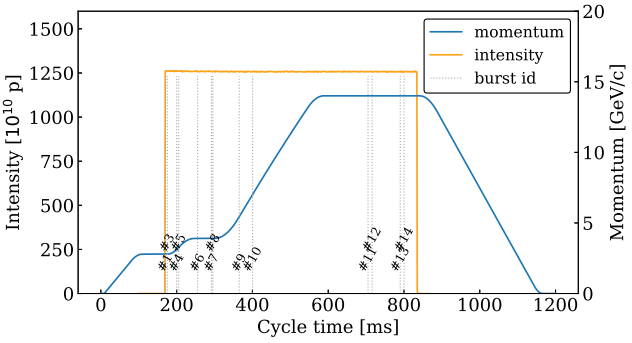

(a) Evolution of the magnetic field and the intensity along the cycle. The vertical gray lines represent the start time of each burst.

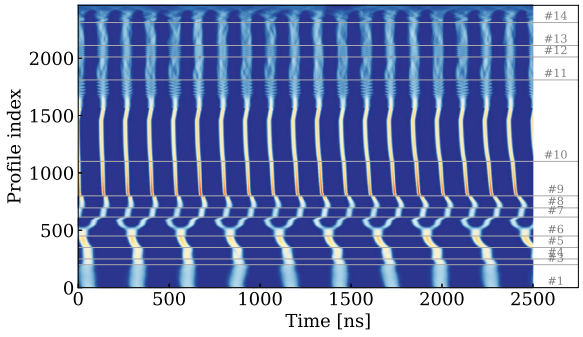

(b) The PS receives eight bunches from the PSB in a single injection, and a double splitting is performed on each bunch before the beam acceleration. At the end of the cycle, the beam is debunched, and a special wideband RF cavity creates a barrier bucket to generate a depleted region to reduce beam loss during the rise of the extraction kickers. Each burst is separated by a gray line in the plot. This figure shows a longitudinal instability that arises around the 2000th profile. The color in the image represents the signal amplitude in arbitrary units.

Fig. 9. Longitudinal beam profile acquisitions on the PS cycle used for fixed targets physics in the SPS North Area. For these acquisitions, the burst times nearly overlap to ensure that the measurements cover the whole cycle. At the end of the cycle, a wideband RF cavity generates a depleted region to reduce beam loss at extraction [21].

3. Longitudinal beam observation acquisitions

The FESA class has been validated across the majority of CERN synchrotrons. This section presents representative operational use cases through its various modes and functionalities. Its initial development was driven by the specific requirements of the PS, which imposes the most significant longitudinal transformations on the beam. Subsequently, the class was generalized to accommodate the diverse operational demands of the other accelerators.

In its present configuration, the PS installation provides up to 16 independently configurable trigger bursts, thereby enabling an equal number of acquisition windows within a single cycle. In the PSB, composed of four superposed rings, one digitizer is used per ring, each supporting up to four bursts. The SPS is equipped with two digitizers: one triggered at each injection, and another offering up to five freely configurable bursts. For the LHC, the new acquisition system has been tested on one full turn profile acquired every second, on each of the two counter-rotating beams.

3.1. PS

To serve the varied user community with beams of significantly different characteristics, a large number of RF systems are installed in the PS. With these RF systems, the longitudinal beam structure can be shaped according to the needs of the user [6] via different RF gymnastics. For example, the beam can be longitudinally split into two or three independent bunches, and a bunch rotation can be performed to shorten the bunch length just before extraction. The PS RF system allows up to 84 bunches with a very short bunch length at the end of the cycle for an LHC-type beam, meeting the SPS longitudinal beam requirements.

Fig. 8 shows the evolution of different bunches along a special LHC-type cycle using the so-called Batch Compression, Merging, and Splitting (BCMS) [22] scheme. The BCMS beam was the operational beam delivered by the injectors to the LHC in 2024–2025. This cycle requires two injections from the PSB into one PS cycle, which can easily be identified on the waterfall plot as the number of bunches increases from three to six from the first to the second burst. The six bunches are then compressed and merged into three bunches. Each bunch is further split into three before acceleration. At the top energy, two double splittings are applied. At the end of the acquisition, the plot shows a beam movement to the right, which is caused by an increase in the revolution period during the rise of the extraction bump. Bunch rotation is applied at the very end, just before the extraction of the beam.

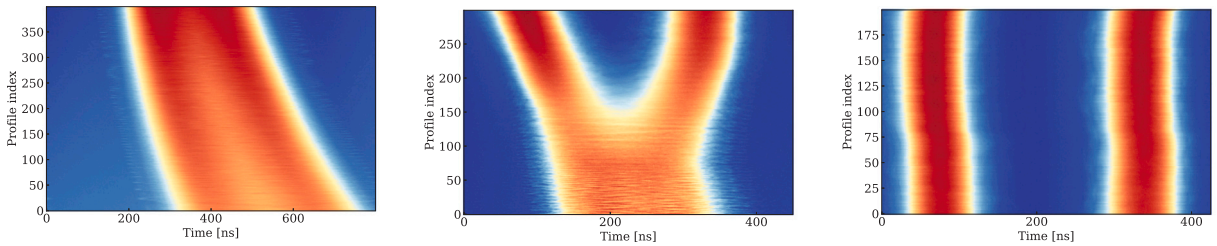

(a) Profiles acquired in one ring of the PSB after the injection from the LINAC of a proton beam.

(b) Profiles acquired in one ring of the PSB during the longitudinal beam splitting.

(c) Profiles acquired in one ring of the PSB at flat top energy of the cycle, just before extraction to the PS

Fig. 10. Waterfall plots of acquisitions at different stages in a PSB cycle using the FESA class burst acquisition mode.

Another acquisition example of a different beam type is shown in Fig. 9. The shown SPS Fixed Target Proton (SFTPRO) beam produced in the PS is specifically designed for the user community at the SPS North Area, and Fig. 9 shows beam instabilities that can be seen in the waterfall plot around the 2000th profile, where bunch-by-bunch variation is visible. This instability was triggered by a longitudinal emittance that was too small during acceleration. This example corresponds to a real operational use case, where the analysis of the waterfall plot allowed the identification of the presence of beam instabilities and the verification of the effectiveness of the mitigation measure (increase of the longitudinal blow-up).

3.2. PSB

The PSB is composed of 4 superposed rings. This design structure was chosen to improve operational flexibility, as each ring is controlled independently, and to increase the available intensity throughput.

In the PSB, the main observation windows occur during injection, end of longitudinal blow-up, and on the flat-top. Depending on the user's requirements, longitudinal bunch splitting might be performed in the PSB to deliver the beam according to the PS specifications. Up to four bursts can be generated for each ring of the PSB, facilitating the observation of these specific moments along the cycle with careful setup of the trigger configuration as shown in Fig. 10.

3.3. SPS

In the SPS, the main points of interest along the proton cycles include measurements after each injection, during the ramp, and on the flat-top corresponding to the highest energy of the cycle. Fig. 11 shows the acquisition of one bunch profile at the first injection into the SPS. At this specific timing, it is crucial to verify that the bunch length falls within the acceptance of the SPS RF parameters to prevent losses. The transfer function of the acquisition system combined with the very short signal amplifies the tail behind the bunch, which requires correction in post-processing (see Section 4.1).

This acquisition system was also used during the 2024 lead ion run in the SPS to observe the momentum slip-stacking of the beam [23]. This complex RF manipulation was developed in the framework of the LHC injectors upgrade project [5] to increase the total number of lead ion bunches injected into the LHC by a factor of two. The slip-stacking process interleaves two batches from separate injections, reducing the bunch spacing from 100 to 50 ns (see Fig. 12).

3.4. LHC

In the LHC, the FESA class can acquire one profile or a set of profiles every second in the continuous mode. This allows the measurement and monitoring of the bunch length evolution over a long time period. Longitudinal tomography can also be performed using a burst of profiles, facilitating the monitoring of the longitudinal emittance.

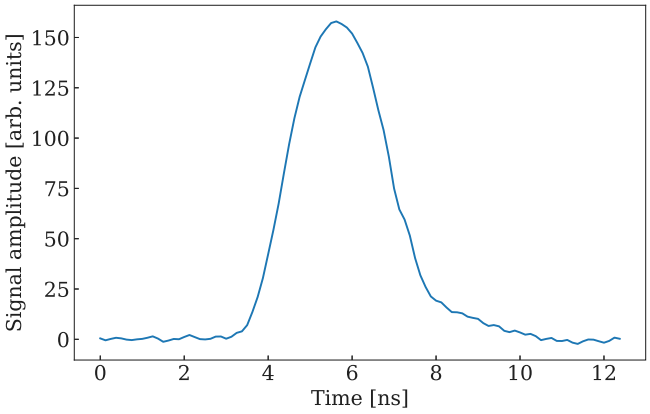

Fig. 11. Raw bunch profile captured just after the injection into the SPS. The transfer function of the acquisition chain causes the profile to be asymmetrical, with a longer bunch tail.

Fig. 13 shows the raw acquisition of a complete turn in one of the two rings in the LHC. The acquisition was performed during stable beam conditions, where the two beams collide in each of the four main experiments in the LHC. During operations, this raw acquisition is taken every second to analyze and derive longitudinal beam parameters, ensuring that the beam observables remain stable over time.

4. Real-time data treatment for longitudinal beam quality observables

A reliable FESA acquisition layer that publishes data for each cycle is the foundation for several downstream applications. These include real-time monitoring of bunch length, bunch-by-bunch intensity, longitudinal emittance via phase space reconstruction, and the optimization of the RF gymnastics. Additional software layers were developed and used specifically for the aforementioned applications. In the following, the focus will be on analyzing and processing raw data to derive longitudinal beam properties from the acquisition data.

While the FESA class is developed in *C++*, most of the code developed for the data analysis runs on a server in *Python*, allowing more flexibility in the libraries that can be used. The Unified Controls Acquisition and Processing (UCAP) [24] framework is a CERN platform for processing and transforming data acquired from any set of devices. It offers processing nodes centralized on a server and a unified way to handle incoming data streams, such as cycle-based grouping of the data published by the devices.

Since most of the developments outlined in the following subsections are computationally intensive, the software analysis layer operates on a dedicated server. In the PS, for example, this includes performing multi-bunch tomography and analyzing approximately 3000

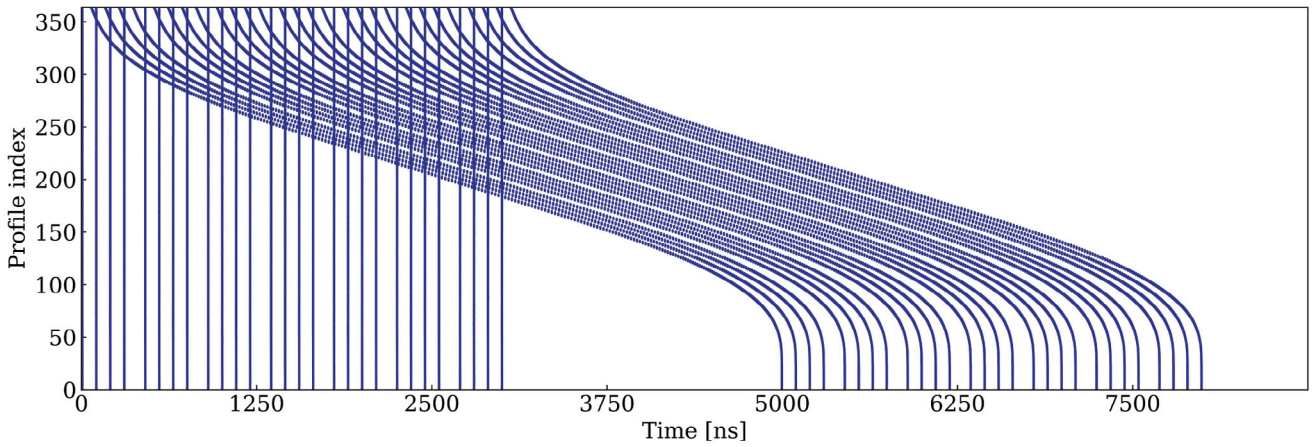

Fig. 12. Momentum slip-stacking of the ion beam in the SPS. Slip-stacking aims to reduce the bunch spacing between consecutive bunches from 100 to 50 ns. The complete slippage of the second batch towards the first one is visible. This process allows for injecting more bunches into the LHC, resulting in a higher luminosity.

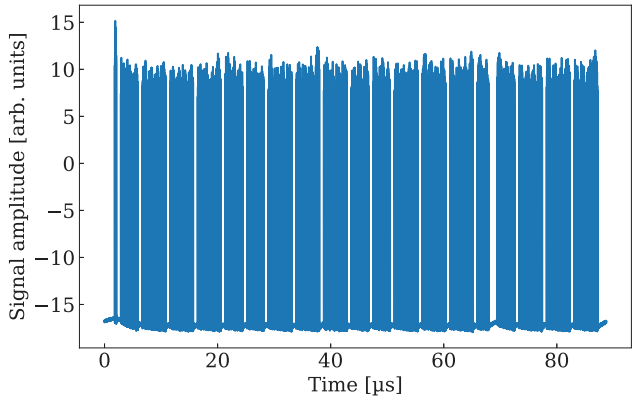

Fig. 13. Acquisition of a complete turn from the LHC Beam 1 during stable beam conditions, when 2460 bunches were colliding in the different experiments.

independent profiles from different bursts on each cycle. The codebase is parallelized, significantly enhancing computation speed and ensuring that analysis results are consistently available before the next cycle data arrival.

4.1. Bunch detection

The first step in the longitudinal beam quality analysis is to efficiently and reliably identify the different bunches from the raw data. Since the real-time monitoring is designed for different synchrotrons, the hyperparameters must be carefully tuned to ensure accurate bunch detection.

The measured raw signal is a convolution of the beam signal with the transfer function of the acquisition chain. The low- and high-frequency parts of the transfer function impact the measured signal in different ways, as shown in Fig. 14. The high-frequency part of the transfer function affects the entire bunch, as the perturbation propagates and accumulates along it. This accumulated effect is particularly significant on the falling edge of the bunch, causing a long bunch tail. This is especially visible for short bunches at PS extraction, and in the SPS and the LHC (see Fig. 14(a)). The low-frequency part of the transfer function causes a baseline droop, resulting in different peak amplitudes along the bunch train if not corrected (see Fig. 14(b)). These effects strictly depend on the acquisition system, and their magnitude varies across accelerators and beam characteristics. The raw profile correction can therefore be divided into two steps: one correcting the high-frequency part and the other correcting the low-frequency part. These two approaches have their own set of assumptions.

The transfer function of the complete acquisition chain can be measured with high precision, ensuring that the phase information necessary to account for the left/right asymmetry is preserved. The resulting effect on the measured signal can then be corrected by applying the inverse transfer function in the frequency domain. By considering *x(t)* as the input signal measured in the WCM and *y(t)* as the output signal after the cables and electronics, the transfer function $H(\omega)$ can be derived as follows:

$$H(\omega) = \frac{\mathcal{F}[y(t)]}{\mathcal{F}[x(t)]}, \tag{1}$$

where $\mathcal{F}$ is the Fourier transform of the system. As expected, the corrected profile has a decreased tail compared to the raw data, as shown in Fig. 14(a). The impact of the transfer function was shown to introduce a discrepancy in bunch length measurement on the order of 15%–20% for very short bunches [25]. This method will correct the high-frequency baseline added by the acquisition system.

The low-frequency part of the transfer function is much more challenging to measure and correct in the frequency domain. Therefore, an alternative approach is employed, which involves the correction in the time domain by considering the inverse effect of an RLC circuit, where C is neglected as its contribution is only significant in the high-frequency part. The parameters of the WCM system in the PS [26] are as follows: 10^{-6} Henries for the inductance (L), 6 Ω for the resistance (R). The sampling interval is defined as Δt, $y(t)$ represents the raw input signal sampled by the digitizer (without correction), and $x(t)$ denotes the output of the digital filter. The low-frequency component of the signal is corrected using the following recursive equation:

$$x\left(t_i\right) = y\left(t_i\right) + \frac{1-K}{1+K} y\left(t_{i-1}\right) + x\left(t_{i-1}\right), \tag{2}$$

where $K = 2L\Delta t/R$ and t_i is the current data sample. This filter linearizes the baseline droop by correcting the RLC effect of the system. Then, the correction of the remaining linear baseline droop can be easily performed between the first and last points of the profile. More details for the low-frequency correction can be found in Appendix: Baseline correction.

The transfer function correction is required to ensure accurate measurement of the longitudinal bunch parameters and constitutes the first step of the proposed algorithm. The baseline is then adjusted such that the zero level corresponds to the minimum value on the *y*-axis. Bunch

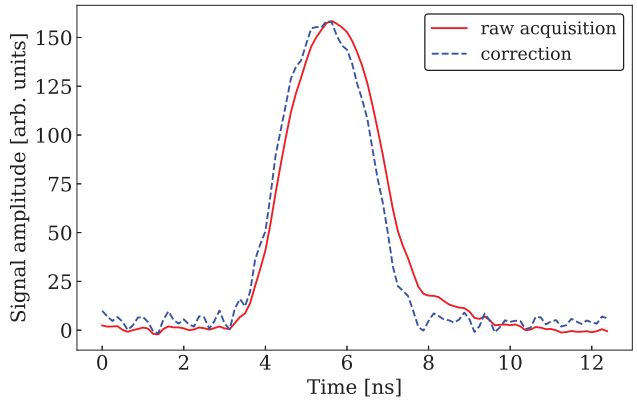

(a) The raw and corrected profiles of one bunch acquired in the SPS. The raw data acquisition in red is affected by the high-frequency part of the transfer function. The profile with the correction of the transfer function is shown in blue. This allows the removal of the tail from the profile. The corrected profile has in comparison more noise in its signal.

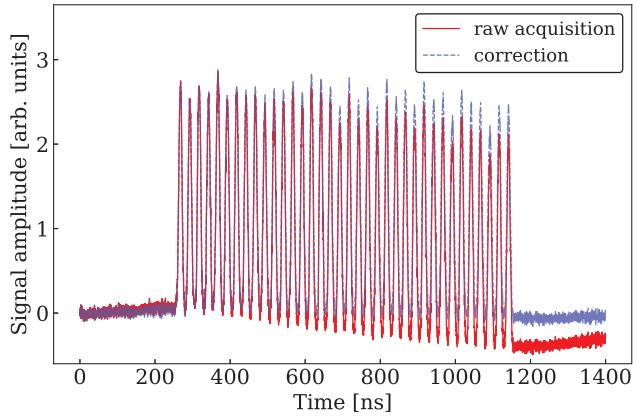

(b) The raw and corrected profiles of an LHC beam acquired in the PS at the end of the cycle. The raw data acquisition in red is affected by the low-frequency part of the acquisition system transfer function. The profile with the correction of the transfer function is shown in blue. By removing the baseline droop, the different bunches become comparable to each other.

Fig. 14. Impact of the transfer function of the acquisition chain on the measured signal.

detection is finally performed on the corrected signal, providing key parameters including the bunch centers and start and end indices.

The bunch detection algorithm uses the *find_peaks* function that can be fine-tuned for this use-case by specifying some parameters such as the *height* (minimum value on the y-axis for a point to qualify as a peak), the *distance* (minimum separation between peaks), and the *width* (required peak width on the x-axis in samples).

This solution strikes a balance between speed and accuracy and is applied to all acquired raw profiles before performing the bunch length analysis and the tomographic reconstructions, as further detailed in Sections 4.2 and 4.3. For instance, as shown in Fig. 8, the algorithm has to process nearly 3000 profiles in less than 1.2 s to publish the results on time.

4.2. Bunch-by-bunch longitudinal analysis

Having completed the bunch detection algorithm for each corrected profile, the bunch length is computed considering the acquisition sampling rate and the number of samples between the left and right limits of each bunch. The "90% bunch length" is then computed from the left and right limits at 10% from the maximum value, and the full width at half maximum (FWHM) length is computed from the limits at 50% of the maximum. However, the standard convention used in operation is to take the 4σ-bunch length as the reference, where σ is the standard deviation of a Gaussian distribution. This 4σ-value is considered as the full bunch length and is therefore tracked throughout the cycle. The computation of the standard deviation for a discrete random variable is given by:

$$\sigma^2 = \sum_{i=1}^{N} p_i \cdot x_i^2 - \left(p_i \cdot x_i\right)^2, \tag{3}$$

where x_i is the sample index and the individual probabilities p_i are defined as

$$p_i = \frac{y_i}{\sum_{i=1}^{N} y_i}, \tag{4}$$

with y_i being the individual signal values.

Bunch spacing is determined by measuring the distance in samples between consecutive bunch centers and multiplying this value by the sampling rate. The total intensity of the beam is measured in units of 10^{10} charges, and is derived from a beam current transformer (BCT [17]) that gives the absolute beam current. To distribute the total intensity to the detected bunches, the area under the bunch curve is normalized per profile and is then multiplied by the total intensity measured from the BCT. It is important to note that this approach is accurate only if all the bunches circulating in the synchrotron are contained within the profile acquisition window.

These metrics can be averaged per profile to identify any trend in their evolution throughout the cycle. Fig. 15 shows the evolution of the mean 4σ bunch length during the acceleration process in the PS. The bunch length decreases until it reaches its minimum at transition crossing, and increases slightly afterwards. All these computed metrics are logged into a database, allowing offline analysis.

4.3. Longitudinal phase space tomography

Knowledge about the longitudinal phase space parameters is especially important in view of assuring the reliability of RF gymnastics and the longitudinal stability of the beam along the cycle in general. Furthermore, the transverse emittance measurement relies on the precise knowledge of the longitudinal momentum spread $\Delta p/p$.

The current technique to obtain information about the longitudinal phase space for an individual bunch is based on the tomographic reconstruction method [27], which reconstructs an n-dimensional object using several $n-1$ dimensional projections. Specifically, the 2D (n) phase-space data are reconstructed from the 1D (n−1) baseline-corrected profiles. Key performance indicators such as $\Delta p/p$, bunch length, longitudinal emittance, and matched area can be computed from the phase space population density.

Until 2024, longitudinal phase space tomography was operationally performed only on demand and on individual bunches. However, knowledge of these observables along the bunch train is especially important for LHC-type beams to evaluate the spread within the bunches and the average performance throughout the year. It allows the operators to act immediately once a performance degradation is observed.

Ref. [28] proposed an automated tomography system for a single bunch, which was developed as a demonstrator for the PSB. In the following, an extension of this demonstrator into an operational system running multi-bunch tomography in real-time for the whole CERN accelerator complex is presented. This extension was eased by the new FESA acquisition class presented in Section 2.

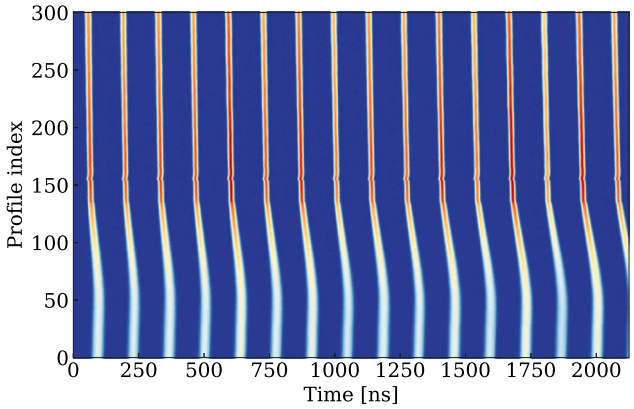

(a) Waterfall plot of one acquisition with 8500 points acquired per profile and 300 profiles taken during the acceleration process. Transition crossing is visible as a slight signal shift around the 150th profile.

(b) Evolution of the mean 4σ bunch length along the different acquisition profiles. As expected, the bunch length decreases in the first part, when the particles are below the transition energy, before increasing afterwards.

Fig. 15. Evolution of the bunch length during acceleration in the PS, around transition crossing.

4.3.1. Longitudinal tomography algorithm

The code implementation of this longitudinal tomography algorithm is detailed in Ref. [29]. It consists of two main stages: particle tracking and an optimization algorithm that reconstructs the longitudinal phase space based on a given profile.

In the tracking stage, particles are represented by their phase and energy coordinates in phase space. This stage computes the evolution of particle density in phase space from the first to the last profile. It highlights the motion of particles within their RF bucket over the acquisition window, including the non-linearity of the synchrotron motion. The algorithm begins with a homogeneous distribution of test particles across the reconstruction region in phase space and generates two arrays representing the x and y coordinates of all particles for each time profile. The resulting density maps serve as initial weights for the reconstruction stage. The tracking stage accounts for the majority of the computational load.

The reconstruction algorithm iteratively updates the density weights in phase space using the measured baseline-corrected profiles. It then projects these updated densities from phase space back into the time domain to generate a reconstructed set of profiles. The discrepancy between the measured and reconstructed profiles is used to refine the phase space density weights in the following iterations.

To reduce execution time for operational applications, the reconstruction algorithm implementation can reuse prior tracking results for configurations where machine settings, such as magnetic field, RF voltage, harmonic number h, etc., remain stable over time. In those cases, initial conditions do not change between reconstructions, and the simulated density maps used as an initial guess for the reconstruction algorithm are still valid. This significantly reduces runtime, enabling real-time multi-bunch tomography for fast-cycling machines, where the algorithm runs in less than one second for tens of bunches.

4.3.2. Real-time accelerator parameters

The tracking algorithm requires various machine parameters as input to simulate particle trajectories in the longitudinal phase space over the acquisition window. The main parameters required for the tomography algorithm are related to geometrical and magnetic properties of the accelerator lattice, the RF cavity voltages and phases, the particle type, the magnetic dipolar field, and additional information specific to the acquired data. All these parameters are both accelerator- and cycle-dependent.

The process of assembling all these parameters varies significantly between accelerators, as the RF systems in the PSB, the PS, and the SPS are completely different. For most beam types, the PSB operates with a single harmonic RF system throughout the cycle, and hence no second harmonic RF system needs to be considered. The PS uses multiple RF systems operating at frequencies from 2.8 to 200 MHz, and their use depends on the beam type. Only one of the systems (C10 system) supports a wide frequency range from 2.8–10 MHz, whereas all other cavities operate with fixed frequency and harmonic values. In the SPS, two cavity groups (200 and 800 MHz) operate with fixed frequencies and harmonic values during the cycle. An additional software layer was developed to collect all relevant accelerator- and cycle-specific parameters to enable the operational deployment of real-time tomographic reconstruction.

4.3.3. Real-time tomography

Several steps are required to prepare the input data and perform automated tomographic reconstruction. The first step is to isolate a single bunch from the acquired data, leveraging the data processing described in Section 4.1. The bunch centers are determined, and the data around the bunch is cropped according to a user-defined percentage of the bucket length. The set of turn-by-turn data profiles must cover at least half a synchrotron period to allow a correct reconstruction of the longitudinal phase space. Furthermore, the information about the accelerator parameters as described in the previous section is required to perform particle tracking under the corresponding conditions.

The first time a cycle configuration is encountered, the tracking results are stored on disk. The tracking results are loaded and reused directly when the same cycle configuration is met again. The second step is the tomographic reconstruction of the longitudinal phase space, which involves iterative projection and back-projection. This workflow is primarily designed for operational setups, where most cycles are repeated without changes to accelerator parameters.

Fig. 16 shows a typical result of the automated tomography, including the waterfall representation of the bunch profiles, the reconstructed longitudinal phase space at a specific profile, and the projections on the x-axis and the y-axis, representing the bunch phase and energy, respectively. The longitudinal parameters derived from the phase space are also shown, such as the emittance, the matched area, the $\Delta p/p$, and the bunch length.

The current setup enables the computation of tomographic reconstructions of the phase space at multiple timings within a cycle. As mentioned in Section 2.1, the PSB has 4 burst windows for which the tomographic reconstruction could be applied, the PS has 16, and the SPS has 5 time windows. All the computed longitudinal metrics are also logged into a database, allowing for the offline analysis of their evolution over time.

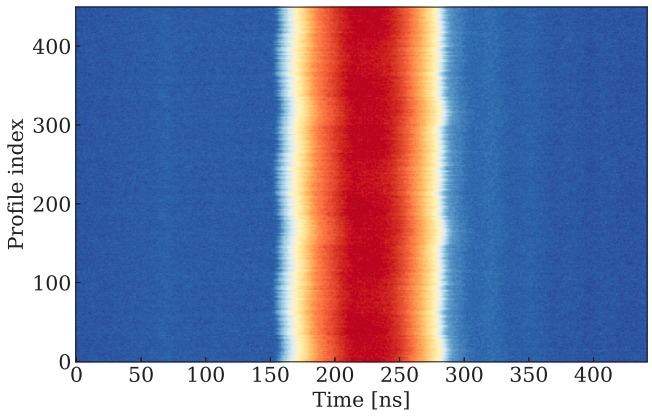

(a) Acquired data of a single bunch from which the tomographic reconstruction is computed.

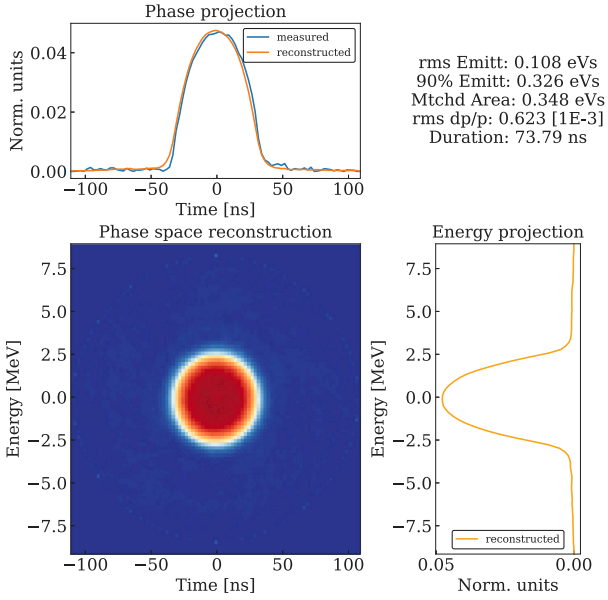

(b) Results from the tomography algorithm. Lower left: reconstruction of the phase space. Top left: reconstructed phase projection from the phase space compared to the acquisition of the specific profile. The reconstruction is in good agreement with the acquisition profile data. Top right: main longitudinal metrics derived from the phase space. Bottom right: reconstructed energy projection from the phase space.

Fig. 16. Typical result of longitudinal tomography before extraction in the PSB.

5. New possibilities in longitudinal beam diagnostics and optimization

This section outlines several advancements enabled by the new longitudinal beam observation framework introduced in this paper. The primary innovation is the unification of the previously very different approaches to longitudinal beam observation and data logging. This common infrastructure allows seamless tracking of beam parameters from one machine to another, ensuring measurement continuity over extended periods. Moreover, the availability of continuous online analysis supports the development of novel optimization algorithms that can dynamically adjust machine parameters based on real-time data.

5.1. Operational GUI application

To simplify configuration and facilitate the visualization of results from the various processing stages described in Section 4, the operational GUI presented in [30] has been further developed to support all new online outputs generated by the longitudinal beam observation framework. The GUI consists of two main components: one dedicated to the configuration of settings and the other to data visualization.

In more detail, the configuration part enables users to define trigger burst parameters, including their timing within the cycle, the number of turns between triggers, and the number of profiles. Additionally, it allows users to specify acquisition parameters per burst, such as the number of samples, the sampling interval, and trigger delays. It also controls the digitizer device settings such as offset, sampling interval, and signal attenuation. Additional panels are available to configure bunch profile analysis and tomography algorithm parameters.

The visualization component features multiple displays to present all online results computed for each cycle. Examples of the application interface are shown in Fig. 17. The GUI provides four primary views: (1) baseline-corrected data from the digitizer, (2) bunch-by-bunch analysis across the whole cycle, (3) a focused bunch-by-bunch analysis for a single profile, and (4) tomography results. These visualizations update automatically as the cycle is executed.

Fig. 17(b) presents the analysis of the final turn in the PS, just before the beam extraction to the SPS. For each burst, a selected profile can be further analyzed using an online processing pipeline on UCAP [31], which replicates the functionality of the existing Bunch Shape Measurement (BSM) [10] operational application. The algorithm finds individual bunches within the profile and applies fitting algorithms to each bunch, such as Gaussian, parabolic, and foot tangent fits. Based on these fits, bunch length and intensity are calculated on a bunch-by-bunch basis. All of the computed results are logged to the database and are made available for inspection through the application, enabling users to observe bunch-by-bunch variations across different time windows within the cycle.

The application is available for each injector, with minor configuration changes that depend on the trigger generator system or the accelerator-specific constraints.

5.2. Bunch splitting automation

The main reason for the versatility of beam production in the PS is based on the various RF systems installed in the accelerator. The different beams can undergo various RF manipulations based on the user's requirements. In the BCMS LHC-type beams, the bunches undergo several bunch splittings. The first splitting divides each bunch into three, and each of these bunches is then split again into four via two double splittings. The triple and two-double splitting procedures are described in detail in Ref. [6]. Different groups of the 10 MHz RF cavities operating at harmonic numbers 7, 14, and 21 ensure the triple splitting process, following a dedicated control program. Similarly, the two-double splittings involve the 20 and 40 MHz cavities to first split the beam from $h = 21$ to $h = 42$, and subsequently from $h = 42$ to $h = 84$ (see also bursts 9 and 13 in Fig. 8).

The performance of those processes drifts over time, leading to variations in bunch-by-bunch intensity and length, deviating from the ideal equal distribution. To counteract the drifts, operators do manual adjustments to the bunch splitting parameters, i.e. cavity phase and voltage, on a regular basis. The newly introduced longitudinal beam observation framework provides systematic and persistent observability of the beam. This is crucial to develop continuous, algorithm-driven optimization of the PS splitting processes. Recent advancements in this area [32] have directly built upon the capabilities offered by this framework. Bunch splitting controllers were built for the triple and double splittings, which continuously monitor and adjust the machine parameters in operations to optimize the beam performance for the LHC-type beams. Their objective functions directly take into account the bunch intensity variation just after the splitting processes and act to ensure equally distributed bunch parameters.

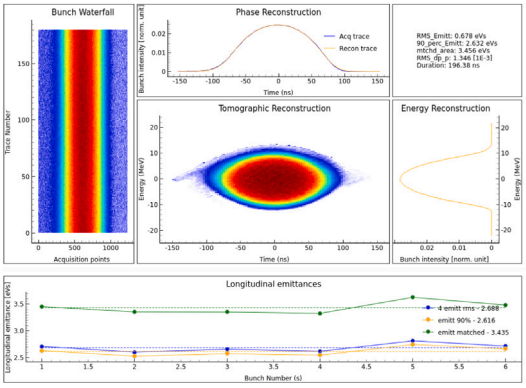

(a) Visualization of the tomographic reconstruction of one out of six bunches just before the triple splitting of an LHC-type beam in the PS. The bunch-by-bunch variations of the emittance and momentum spread over all six bunches are also available.

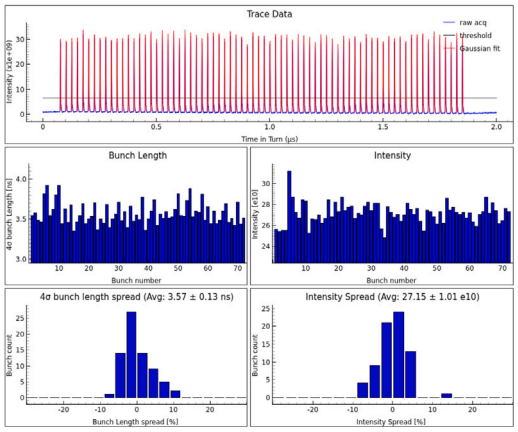

(b) Visualization of the last profile of the cycle. The bunch-by-bunch variations of the bunch length and the bunch intensity of the beam sent to the SPS are shown.

Fig. 17. Different graphical representations of the data available in the operational application.

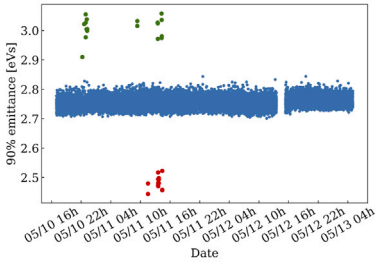

(a) Each point represents the "90% emittance" measurement of the beam at the flat top of the cycle.

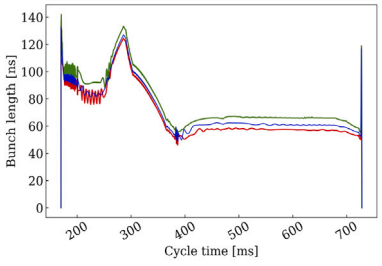

(b) Evolution of the bunch length along the machine cycle. The color-coding highlights the correlation with the outliers seen in Fig. 18(a). The bunch length has bigger oscillations for the green and red curves at the beginning of the cycles.

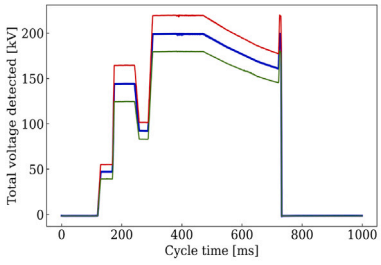

(c) Evolution of the mean total voltage detected along the cycle in the machine. The color-coding highlights the correlation with the outliers seen in Fig. 18(a).

Fig. 18. Evolution of longitudinal parameters for the PS beam to the nTOF facility over three days correlated with the total measured RF cavity voltage. The abnormal points are highlighted in red for the lower emittance values and in green for the higher ones. This analysis directly correlates a difference in longitudinal emittance with a variation in bunch length that originates from the cavities, which were erratically pulsing at an incorrect voltage, letting the bunch have quadrupolar oscillations.

5.3. Beam performance tracking

Since the new framework logs every computed result for all produced beams to a database, it enables the collection of longitudinal beam parameters over an extended time window that can be used for analysis and evaluation. This allows to identify slow parameter drifts and erratic shots, and now offers the possibility of investigating their correlation with other machine signals.

An example of one study is presented in Fig. 18, which shows data from the TOF cycle in the PS over a three-day period. The analysis focuses on the evolution of the "90% emittance" value for all beams generated by the PS and delivered to the neutron Time-Of-Flight (nTOF) facility [33]. The "90% emittance" is defined as the longitudinal phase-space area enclosing 90% of the particles. Fig. 18(a) indicates that the longitudinal beam emittance is generally stable, typically ranging between 2.7 and 2.8 eVs. However, some outlier shots deviate significantly from this range, as indicated by the green and red dots in the figure.

By cross-referencing these deviations with bunch length measurements, a systematic correlation between a variation of emittance and bunch length was identified. Further examination of bunch length evolution across all profiles in the cycle revealed that significant quadrupolar oscillations are visible in the bunch length parameter for these cycles compared to a typical scenario, as shown in Fig. 18(b). This behavior indicated that the RF cavities were incorrectly pulsed during these specific cycles. This conclusion was verified, and the results are shown in Fig. 18(c), where the total voltage of the RF cavities is shown over the complete cycle. It is clearly visible that erratically pulsing RF cavities can explain the outliers in terms of longitudinal emittance and bunch length.

This simple example illustrates how variations in beam parameters can be directly correlated with the performance of the equipment installed in the accelerator. The presence of the new longitudinal beam observation framework enables completely new ways of investigating the performance of the accelerator and identifying the root cause of issues.

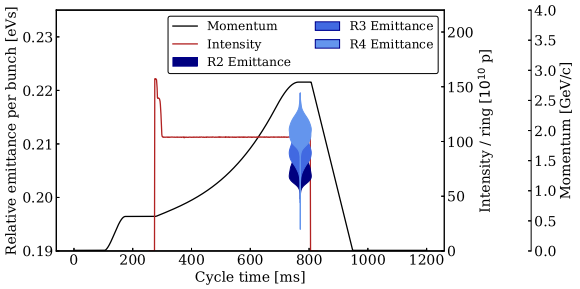

(a) Evolution in the PSB. Each violin plot represents the relative emittance measurement per ring of the beam at the flat top energy of the cycle, just before its extraction.

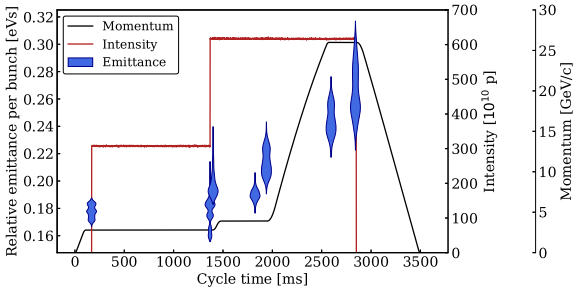

(b) Evolution in the PS. The violin plots show the first and second injection from the PSB, the start of the intermediate plateau, before and after the triple splitting, the beginning of the flat top energy, after the two double splittings, and before the extraction of the beam to the SPS. Each of the splittings introduces an uncontrolled longitudinal blow-up to the beam.

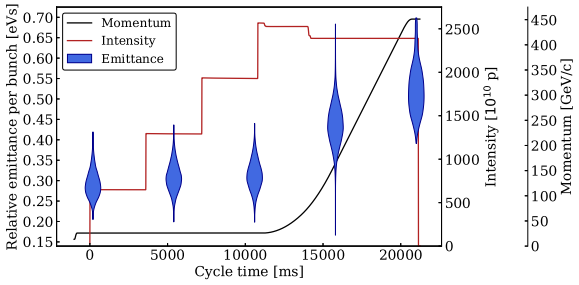

(c) Evolution in the SPS. The first three violin plots show the normalized emittance after the first three injections. The two remaining ones are taken during the beam acceleration and at the flat-top of the SPS. The SPS introduces a controlled longitudinal blow-up during the acceleration to increase the bunch stability. Since no further splitting is performed, the splitting ratio is one, and the normalized emittance equals the emittance.

Fig. 19. Evolution of the normalized emittance per bunch across the different injector synchrotrons. Normalization is performed by taking into account the remaining splitting ratio per bunch. The intensity and momentum evolution are displayed to ease the identification of the different cycle times.

5.4. Across the accelerator chain analysis

The beams can be measured at different moments in each accelerator, and correlation with other parameters, such as transmission of beam intensity, is now possible. For each splitting, the longitudinal emittance is divided by the number of resulting bunches. By convention, we use the normalized emittance after all splittings to compare longitudinal emittances. To compute the *normalized emittance*, each bunch emittance is scaled by the number of remaining splittings it must undergo before reaching its final state. For example, after the triple splitting in the PS, the beam still has a remaining splitting factor of 4, since the subsequent double splittings will further divide the bunches longitudinally. By comparing the normalized emittance, the evolution of beam performance can be consistently tracked across accelerators, potentially highlighting sources of degradation. Fig. 19 shows the evolution of the normalized emittance per bunch in the different injectors during an LHC fill.

The operational beam taken by the LHC is the batch compression, merging and splitting (BCMS) beam [22], which requires two injections of three bunches from three independent rings in the PSB to produce a total of 36 bunches at PS extraction. Fig. 19(a) shows the distribution of the normalized emittance for each PSB ring, just before the extraction to the PS. Then the performance evolution is shown in the PS in Fig. 19(b), where the first two violin plots show both injections from the PSB. Subsequently, the emittance is computed at the beginning of the intermediate plateau, before the batch compression. The triple splitting happens just before the beam acceleration (around 1830 ms), and the normalized emittance per bunch is computed before and after the splitting. The normalized emittance is again computed at the flat top energy, before and after the two double splittings. In the SPS, the same analysis is shown in Fig. 19(c). The emittance is computed after the first three injections. The beam is scraped in the transverse plane during acceleration, shown by the beam loss before the acceleration (around 14 s). To ensure the stability of the bunches during acceleration, the beam is longitudinally blown up, and this effect is visible in the last two violin plots.

In a similar manner to the evolution of the normalized emittance, any other longitudinal parameter can also be tracked using the new framework.

6. Conclusion

A new real-time monitoring framework for longitudinal beam quality has been developed to address the limitations of the existing acquisition and analysis software at CERN. The framework combines a low-level acquisition layer with dedicated server-based data-processing layers, providing a unified approach to data acquisition and analysis. It supports continuous, cycle-by-cycle acquisition and is highly configurable to be deployed at any synchrotron at CERN. An operational application manages configuration settings and allows the visualization of processing results across the framework layers. This enables real-time derivation of key metrics, such as bunch length and longitudinal emittance, for continuous beam quality monitoring during operation.

In addition, the acquisition class is modular and extensible to multi-channel digitizers, such as the head-tail monitor pick-up, providing a groundwork for broader applications. This system also enables the implementation of real-time optimization layers, allowing autonomous adjustment of beam parameters based on live feedback. An example of these optimizers is already operational in the PS, where it optimizes longitudinal bunch splitting to enhance beam quality and streamline LHC delivery.

The system is fully integrated into daily operations, supporting real-time correlation of longitudinal properties with other machine parameters and enabling active optimization. This represents a significant advancement in the longitudinal diagnostics toolkit and paves the way for future developments in autonomous beam optimization across the CERN accelerator complex.

CRediT authorship contribution statement

Amaury Beeckman: Writing – review & editing, Writing – original draft, Validation, Software, Methodology, Investigation, Conceptualization. **Alexander Huschauer:** Writing – review & editing, Writing – original draft, Supervision, Methodology, Conceptualization. **Alexandre Lasheen:** Writing – review & editing, Supervision, Methodology, Conceptualization. **Cédric Hernalsteens:** Writing – review & editing, Supervision. **Nicolas Pauly:** Writing – review & editing, Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors would like to thank T. Argyropoulos, S. Albright, G. Trad, and M. Cejp for their support and the many fruitful discussions.

Appendix. Baseline correction

This appendix explains the assumptions used to derive the equations to correct the baseline droop, which is affected by the acquisition system transfer function. A Wall Current Monitor can be modeled as a parallel RLC circuit with the following impedance.

$$H(\omega) = \frac{1}{\frac{1}{R} + \frac{1}{j\omega L} + j\omega C}$$

The RLC model, together with the measured transfer function of the PS WCM [26], is shown in Fig. 20. The same analysis can be done for the SPS and LHC WCM [34]. The parameters used to model the PS WCM transfer function are as follows: the resistance (R) is 6 Ω, the inductance (L) is 10^{-6} H, and the capacitance (C) is 6.5×10^{-12} F. The cutoff frequencies based on the RLC parameters are given as f_L equals 95.5 kHz and f_C is 4.1 kHz

$$f_L = \frac{\omega_L}{2\pi} = \frac{R}{2\pi L}$$

$$f_C = \frac{\omega_C}{2\pi} = \frac{1}{2\pi RC}$$

The treatment of the transfer function at high frequency can easily be done with Fast Fourier Transform with sufficiently high sampling rate [25]. Correcting the high-frequency part of the transfer function will affect the bunch-profile component at frequencies above the beam spectrum, on the order of 10^8 Hz and higher. The perturbation due to the roll-off at high frequency is then only relevant for short bunches, in the ns scale. The treatment of the baseline droop due to the low-frequency cut-off is not straightforward in the frequency domain, as it requires a very long acquisition to resolve below the revolution period, which is not typical for longitudinal beam observation acquisitions. The

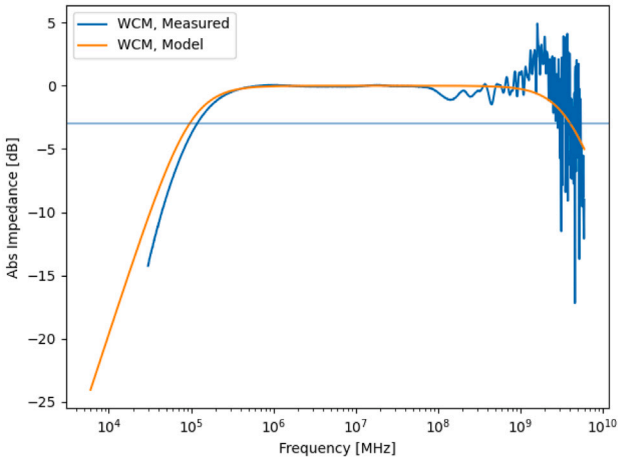

Fig. 20. The transfer function of the PS WCM, as measured and applying the RLC circuit model. f_L and f_C correspond to the 3 dB cut-off, respectively at low and high frequencies.

treatment of the low-frequency part of the transfer function will correct the impact on the profile for frequencies below 10^6 Hz.

It is then assumed that both types of input-signal degradation occur in sufficiently distinct frequency regimes to be treated separately. The expression can be simplified to focus on low frequencies, hence neglecting the capacitance.

$$H(\omega \ll \omega_C) = \frac{j\omega\tau}{1 + j\omega\tau}$$

where $\tau = 1/\omega_L = L/R$ is the time constant. The impedance is turned into a digital filter using the bilinear transform with the following substitution

$$j\omega = \frac{2}{T_s}\frac{1 - z^{-1}}{1 + z^{-1}}$$

where z is the complex frequency from the discrete Laplace transform and T_s is the sampling interval. Deriving through lead to the following transfer function of an Infinite Impulse Response (IIR) filter

$$H(z) = \frac{b_0 + b_1 z^{-1}}{a_0 + a_1 z^{-1}}$$

where

$$a = \left(1, \frac{1 - K}{1 + K}\right),\ b = \left(\frac{K}{1 + K}, \frac{-K}{1 + K}\right)$$

and $K = 2\tau/T_s$. The inverse filter can be obtained by exchanging the a and b coefficients (inverted coefficients are now noted with i subscript). After re-normalizing to impose $a_{i,0} = 1$, the expression can be rewritten as

$$a_i = (1, -1),\ b_i = \left(1, \frac{1 - K}{1 + K}\right).$$

The corresponding time domain digital filter to invert the baseline due to the low frequency cutoff is then

$$x[t_i] = y[t_i] + \frac{1 - K}{1 + K} y[t_{i-1}] + x[t_{i-1}]$$

where x is the reconstructed input sample to the WCM and y is the measured output sample. It is to be noted that the filter will not correct DC offsets and can introduce a linear baseline to the acquisition frame if the DC offset is not zero. This can be easily treated by applying a linear baseline correction, assuming that the first and last samples of the acquisition correspond to samples with no beam.

Data availability

Data will be made available on request.

References

[1] R. Catherall, W. Andreazza, M. Breitenfeldt, et al., The ISOLDE facility, J. Phys. G: Nucl. Part. Phys. 44 (9) (2017).

[2] P. Burdelski, et al., CERN Proton Synchrotron East Area Facility: Upgrades and Renovation During Long Shutdown 2, Number CERN-2021004 in CERN Yellow Reports: Monographs, CERN, Geneva, Switzerland, 2021.

[3] D. Banerjee, J. Bernhard, M. Brugger, et al., The North Experimental Area at the CERN Super Proton Synchrotron, Technical Report CERN-ACC-NOTE-2021-0015, CERN, Geneva, Switzerland, 2021.

[4] S. Deghaye, E. Fortescue-Beck, Introduction to the BE-CO Control System, Number CERN-ACC-NOTE-2020-0069, CERN, Geneva, Switzerland, 2020, URL https://cds.cern.ch/record/2748122.

[5] H. Damerau, A. Funken, R. Garoby, et al., LHC Injectors Upgrade: Technical Design Report, Volume 1: Protons, Number CERN-ACC-20140337, CERN, Geneva, Switzerland, 2014.

[6] S. Gilardoni, D. Manglunki, Fifty Years of the CERN Proton Synchrotron: Volume 1, Number CERN-2011-004 in CERN Yellow Reports: Monographs, CERN, Geneva, Switzerland, 2011.

[7] P. Forck, Juas ii.9 – beam instrumentation, in: CERN Yellow Reports: School Proceedings, vol. 2024–003, 2024, pp. 1339–1528.

[8] S. Deghaye, D. Jacquet, I. Kozsar, et al., OASIS: A new system to acquire and display the analog signals for LHC, in: Proc. ICALEPCS'03, JACoW Publishing, Gyeongju, Korea, 2003, pp. 359–361, number WP502.

[9] J.-F. Comblin, S. Hancock, A. Sanchez, et al., A Pedestrian Guide to Online Phase Space Tomography in the CERN PS Complex, Technical Report CERN-PS-RF-Note-2001-010-Rev, CERN, Geneva, Switzerland, 2003, URL https://cds.cern.ch/record/2840544.

[10] J.L. Alvarez, M. Ruette, PS Bunch Shape Measurement, Technical Report CERN-PS-OP-Note-98-09-Tech, CERN, Geneva, Switzerland, 1998.

[11] A. Guerrero, J.J. Gras, J.-L. Nougaret, et al., CERN front-end software architecture for accelerator controls, in: Proc. ICALEPCS'03, JACoW Publishing, Gyeongju, Korea, 2003, pp. 342–344, number WE612.

[12] G. Papotti, A beam quality monitor for LHC beams in the SPS, in: Proc. EPAC'08, JACoW Publishing, Genoa, Italy, 2008, pp. 3324–3326, number THPC144.

[13] T. Hoffmann, FESA – the front-end software architecture at FAIR, in: Proc. PCaPAC'08, JACoW Publishing, Ljubljana, Slovenia, 2008, pp. 183–185, number WEP007.

[14] C. Roderick, R. Billen, The LSA Database to Drive the Accelerator Settings, Technical Report CERN-ATS-2009-100, CERN, Geneva, Switzerland, 2009.

[15] M. Saccani, E. Calvo Giraldo, W. Viganò, et al., The beam loss monitoring system after LHC injectors upgrade at CERN, in: Proc. IBIC'21, JACoW Publishing, Pohang, Korea, 2021, pp. 285–289, number TUPP32.

[16] G. Valentino, G. Baud, R. Bruce, et al., Final implementation, commissioning, and performance of embedded collimator beam position monitors in the large hadron collider, Phys. Rev. Accel. Beams 20 (8) (2017).

[17] M. Andersen, Fast beam current transformer software for the CERN injector complex, in: Proc. ICALEPCS'11, JACoW Publishing, Grenoble, France, 2011, pp. 382–385, number MOPMS027.

[18] Acqiris, Acqiris SA217P 14-bit ADC Card — Datasheet, Technical report, Acqiris, 2021.

[19] Acqiris, Acqiris SA248P 14-bit ADC Card — Datasheet, Technical report, Acqiris, 2021.

[20] F. Hoguin, M. Jaussi, CSS: OASIS & COHAL, 2024, Electronics-forum #10 joint with Front-End Developers: HW Abstraction Layers.

[21] M. Vadai, H. Damerau, M. Giovannozzi, et al., Implementation of Synchronised PS-SPS Transfer with Barrier Buckets, Technical Report LLRF2022/33, CERN, Geneva, Switzerland, 2022.

[22] S. Albright, M. Bartosik, T. Bohl, et al., Overview of the beams from the injectors, in: 8th Evian Workshop on LHC Beam Operation, CERN, Evian, France, 2017, pp. 119–124.

[23] T. Argyropoulos, T. Bohl, A. Lasheen, et al., Momentum slip-stacking in CERN SPS for the ion beams, in: Proc. IPAC'19, JACoW Publishing, Melbourne, Australia, 2019, pp. 3184–3187, number WEPTS039.

[24] L. Cseppentő, M. Büttner, UCAP: A framework for accelerator controls data processing @ CERN, in: Proc. ICALEPCS'21, JACoW Publishing, Shanghai, China, 2021, pp. 230–235, number MOPV039.

[25] J. Flowerdew, L. Intelisano, A. Lasheen, et al., Cable Transfer Function Measurements and Bunch Length Analysis in the SPS, Technical Report CERN-ACC-NOTE-2025-0018, CERN, Geneva, Switzerland, 2025.

[26] J. Belleman, W. Andreazza, A. Nosych, A new wall current monitor for the CERN proton synchrotron, in: Proc. IBIC'16, JACoW Publishing, Barcelona, Spain, 2016, pp. 143–146, number MOPG41.

[27] S. Hancock, A Simple Algorithm for Longitudinal Phase Space Tomography, Technical Report CERN-PS-RF-Note-97-06, CERN, Geneva, Switzerland, 1997.

[28] S. Albright, C. Grindheim, A. Lasheen, et al., Recent developments in longitudinal phase space tomography, in: Proc. IPAC'22, JACoW Publishing, Bangkok, Thailand, 2022, pp. 347–350, number MOPOPT043.

[29] C. Grindheim, S. Albright, Longitudinal Phase Space Tomography Version 3, Technical Report CERN-ACC-NOTE-2021-000, CERN, Geneva, Switzerland, 2021.

[30] D. Krefta, A. Huschauer, A. Lasheen, et al., Multi-bunch and Multi-burst Longitudinal Beam Observation in the PS, Technical Report CERN-ACC-NOTE-2024-0015, CERN, Geneva, Switzerland, 2024.

[31] M. Coly, Private communication, 2025.

[32] J. Wulff, Optimizing RF triple splittings with reinforcement learning, 2025, Inverted CERN School of Computing.

[33] G. Tagliente, O. Aberle, et al., The n_TOF facility at CERN, EPJ Web Conf. 297 (2024) 01013.

[34] T. Bohl, J.F. Malo, The APWL Wideband Wall Current Monitor, Technical Report CERN-BE-2009-006, CERN, Geneva, Switzerland, 2009.