Automatic detection of transverse beam instabilities in the Large Hadron Collider

T. E. Levens^(D),^{*} K. Łasocha,[†] T. Lefevre, M. Gąsior, R. Jones, T. Włostowski, J. P. Ellis,[‡] and R. J. Steinhagen[§] *CERN, CH-1211 Geneva 23, Switzerland*

(Received 12 June 2019; published 20 November 2019)

Transverse beam instabilities are a major cause of emittance dilution and beam losses in high intensity hadron accelerators. This paper presents an overview of the system that has been developed for automatic detection of transverse beam instabilities in the Large Hadron Collider at CERN. It relies on the detection of very small transverse beam oscillations using a high sensitivity tune measurement system. These oscillations are then processed online in a field programmable gate array to monitor the growth of any instabilities. When an instability is detected, the acquisition of intrabunch beam position data with a multigigahertz bandwidth is triggered. This data allows identification of unstable bunches within bunch trains and provides information on the temporal signature of the instability. The design of these systems and the different algorithms for instability detection that have been tested with beam will be presented along with a detailed description of their performance.

DOI: 10.1103/PhysRevAccelBeams.22.112803

I. INTRODUCTION

For applications requiring high brilliance [1,2] or high luminosity [3] beams, accelerators are being pushed to operate with an ever increasing number of particles. Collective effects resulting from space charge [4], wakefields [5] and electron cloud [6-9] may lead to beam instabilities that cause emittance dilution and beam loss. An exhaustive review of beam instability mechanisms and their possible mitigations are presented in [10]. Detection of beam instabilities is crucial to understand and optimize the operation of such high intensity accelerators. Depending on the beam type and its characteristics, such as bunch length, appropriate beam diagnostics have been developed to measure instabilities in both the transverse and longitudinal planes. For the bunch lengths typically found in hadron accelerators, electromagnetic pickups have been used successfully for many years [11] to allow the detection of beam instabilities with bandwidths of up to a few gigahertz. Transverse instabilities of short electron bunches in storage rings are usually observed using synchrotron radiation [12] and streak cameras [13].

Transverse beam instabilities have been observed in the Large Hadron Collider (LHC) at CERN since the first energy ramp with a single nominal bunch of $\sim 10^{11}$ protons in 2010 [14]. Throughout the first physics run (LHC run 1, 2009–2013), instabilities were regularly observed at injection and during physics fills with 50 ns bunch spacing at an energy of 3.5 TeV [15].

For the second physics run (LHC run 2, 2015–2018), a number of changes were made to the machine operating conditions to increase the luminosity provided to the experiments. The bunch spacing was decreased to the design value of 25 ns, enhancing the influence of electron cloud and other collective effects [16]. Tighter collimator settings, with an associated increase in impedance, was required for running at lower than nominal β^* values [17]. In addition, strict limits on beam loss at the increased operating energy of 6.5 TeV were required to avoid magnet quenches [18]. All of these factors meant that the mitigation of beam instabilities continued to be an important consideration for regular machine operation during run 2.

Beam instrumentation capable of characterizing instabilities is used extensively during machine development sessions to understand the cause of the observed instabilities. A number of studies have been carried out to experimentally qualify the LHC impedance model [19] and optimize machine parameters in order to mitigate instabilities that occur during regular operation. Although this instrumentation was already available during run 1, it required an operator to trigger the acquisition when an

^{*}Also at University of Glasgow, Glasgow, Scotland.

[†]Also at Institute of Physics, Jagiellonian University, Kraków, Poland.

[‡]Also at University of Melbourne, Melbourne, Australia.

[§]Present address: GSI Helmholtzzentrum für Schwerionenforschung, Darmstadt, Germany.

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

instability occurred. This manual procedure was error prone and often resulted in missing the instability if an instrument was triggered either too early or too late. As instabilities can occur at unpredictable times during the machine cycle, it became clear that it would be beneficial to have a system to automatically detect the onset of an instability and trigger other instruments synchronously.

Since the beginning of LHC run 2, a number of improvements to the beam instrumentation were implemented and an "instability trigger network" was developed to connect to these instruments and provide synchronous acquisitions.

II. LHC INSTABILITY TRIGGER NETWORK

The LHC instability trigger network (LIST) enables bidirectional trigger distribution between instruments capable of detecting and observing beam instabilities. If any instrument detects an instability, it can synchronously trigger the acquisition of other instruments via the network to provide a complete snapshot of the beam conditions around the moment of the instability.

The LHC is served by two distinct timing networks, a millisecond synchronous "general machine timing" used by all systems and a turn synchronous "beam synchronous timing" system dedicated to beam instrumentation [20]. While the option to distribute triggers using these existing networks was considered, both networks are unidirectional and are designed to distribute centrally generated timing events to remote equipment. Therefore, neither offers the possibility to receive and redistribute triggers generated by remote equipment without dedicated cabling from the equipment to the network "master" located in the CERN Control Centre. Furthermore, both networks have limited bandwidth available to distribute events, making the addition of many additional triggers to these networks impractical. Because of these limitations, the LIST network was deployed and is dedicated to distribution of instability trigger events.

The LIST network is based on White Rabbit [21] technology, a deterministic, synchronous extension to the Ethernet standard. White Rabbit allows geographically separated nodes on a network to be synchronized with subnanosecond precision. The basic operation of the LIST network is shown in Fig. 1. A trigger pulse generated by an instrument has its arrival time tagged by a time-to-digital converter at an input node of the network. The time stamp of this event is then distributed over the network to all other nodes. As the White Rabbit network provides a synchronized clock at all of its nodes, they can share a common notion of time. The pulse can therefore be regenerated synchronously at any output with a precise delay that can be programmed down to a minimum value defined by the



FIG. 1. A functional diagram of the LHC instability trigger network. At an input node, the arrival time of a trigger pulse is measured with a time to digital converter (TDC). The time stamp is distributed over a White Rabbit network that also provides clock synchronization to the nodes and distributes International Atomic Time (TAI) as a common time base. An output node can regenerate the trigger after a defined delay using a programmable pulse generator.

routing and distribution delays inherent to the network. The design of the LIST is described in more detail in [22].

III. LHC HEAD-TAIL MONITOR

A. Hardware description

An important instrument for characterizing beam instabilities in the LHC is the "head-tail monitor." The system, shown as a simplified block diagram in Fig. 2, is based on the high speed acquisition of a long "stripline" type beamposition monitor (BPM). The LHC head-tail monitor has two dedicated single-plane BPMs per beam, one horizontal



FIG. 2. The principle of the LHC head-tail monitor. The sum (Σ) and difference (Δ) of the electrode signals of a long "stripline" type BPM are obtained using a 180° hybrid and sampled by a high speed digitizer located in a service gallery.



FIG. 3. Time domain signal and corresponding frequency spectrum from a 40 cm stripline BPM for a nominal LHC bunch passage. Each plot shows the measured sum signal (solid line) compared to a simulated bunch approximated with a raised cosine bunch shape (dotted line).

and one vertical, installed at locations with high β -functions in the measurement plane. A commercial wideband 180° hybrid produces the analog sum (Σ) and difference (Δ) of each pair of electrode signals. The signals are directly digitized with a high-speed digitizer located close to the beam line in a service gallery. A head-tail monitor was initially installed in the CERN Super Proton Synchrotron (SPS) to allow chromaticity measurements through the observation of the periodic dephasing between the head and tail of the bunch after a transverse kick is applied to the beam [23,24]. The LHC system was also installed for this purpose, however the emittance blowup resulting from the required kick excitation has meant that it has never been used operationally for chromaticity measurements. Instead, the system has been primarily used for direct time-domain measurements of intrabunch motion, which is possible due to the large bandwidth of the BPM and the acquisition system.

Due to the geometry of the pickup, the impulse response of a stripline BPM has two pulses with opposite polarity that are separated by twice the length of the electrode [25]. In the frequency domain, this results in periodic notches in the response at multiples of twice the electrode propagation delay, where the signals interfere destructively. A measurement of the signal from a single bunch passage through a 40 cm stripline BPM (the type used in the LHC for the head-tail monitor) is shown in Fig. 3(a). The notches in the frequency response occur at multiples of 375 MHz corresponding to $c/(2 \times 0.4 \text{ m})$. For this BPM the electrodes have been dimensioned such that, for the typical LHC bunch length of 1.2 ns (4σ) and 25 ns bunch spacing, the



FIG. 4. S-Parameter transmission measurement from the upstream to downstream port of a 40 cm stripline BPM.

two pulses are well separated in time and do not interfere with the signal from neighboring bunches. This allows the reflected pulse to be removed by gating the signal in the time domain, effectively removing the notches in the frequency response as shown in Fig. 3(b). After gating, the BPM has a usable frequency response up to a few gigahertz, above which imperfections in the BPM geometry and cabling become significant, as shown in Fig. 4.

Although the head-tail monitor can provide a direct measurement of intrabunch motion, it suffers from a number of limitations. The sensitivity of the system, defining the minimum detectable oscillation amplitude, is limited by the available dynamic range of the high-speed digitizers used to acquire the signals. As a consequence of this, the beam orbit must be kept centered in the BPM so that the beam position signal is minimized and the maximum dynamic range can be utilized for digitizing the small oscillation signal. However, even for a perfectly centered beam, imperfections in the hybrid couplers used for the difference calculation result in some residual signal remaining.

A further limitation comes from the amount of acquisition memory available and the speed at which data can be read into the control system. The initial LHC head-tail monitor used commercial 8-bit, 10 GSPS oscilloscopes that were limited to a maximum of 11 turns for all bunches (1 ms of data). This therefore implied the need for precise triggering once the oscillation amplitude had reached a sufficient level to be visible, but before significant beam losses led to a beam dump. Each acquisition took approximately 10 seconds to read out which limited the time between triggers.

In the decade since the LHC start-up, a number of new digitizers have become available with higher performance in terms of analog bandwidth, resolution, and depth of acquisition memory. New 10-bit, 10 GSPS oscilloscopes with 4 GHz analog bandwidth were tested in the SPS during 2017. A comparison measurement is presented in Fig. 5 showing that the new digitizers have significantly better performance resulting in a 9 dB lower noise floor for the same signal. They also offer a total acquisition length of 450 turns (41 ms) and allow much faster data readout, relaxing slightly the requirements on triggering. The new model has been used in the LHC for operation since the 2018 run with excellent results.



FIG. 5. Head-tail monitor acquisition from the Super Proton Synchrotron comparing the old 8-bit oscilloscope (top) and new 10-bit oscilloscope (bottom). Each plot overlays the first 43 turns after injection of a single bunch where, in addition to the injection oscillation, longitudinal and intrabunch motion are visible.

B. Data processing

The instability mode number can be determined visually from the head-tail monitor data by overlapping multiple turns of a single bunch and counting the number of stationary nodes in the oscillation. In order to achieve this the raw data obtained from the head-tail monitor requires a number of postprocessing steps to correct imperfections in the acquisition chain and obtain a "clean" signal. During LHC run 1, as the head-tail monitor was only used for specific measurement campaigns, these processing steps were performed manually. Since the number of acquisitions taken during a study is relatively low, the manual processing was not overly time consuming. However, in order to better understand the instabilities that occur during normal operation, the head-tail monitor is now triggered automatically via the LIST network every time an instability is detected and generates large quantities of data during daily operation. As manual processing of these large data sets is not feasible, automatic algorithms have been developed to perform the necessary postprocessing steps and determine whether an acquisition contains an instability.

The first step is to determine which bunch slots are filled with beam to avoid having to process empty slots. A single turn (89 μ s) of sum signal data is divided into 3564 25 ns intervals each corresponding to approximately one bunch slot. Each of these slots is then subdivided into five 5 ns segments and the signal amplitude in each segment is calculated. For an empty slot, all of the segments will only contain noise and will therefore have similar amplitude. For a slot where a bunch is present at least one of the segments will contain data of a much higher amplitude. The ratio between the highest and lowest amplitude segments can therefore be used to determine the presence of a bunch, with this comparison robust against any changes in the baseline noise level of the signal. The threshold has been set empirically such that "nominal intensity" bunches $(N_b = 1.1 \times 10^{11})$ will be reliably detected while "pilot" bunches $(N_b = 5 \times 10^9)$ will be ignored. These settings avoid saving acquisitions which only contain low intensity bunches, as the head-tail monitor does not have enough dynamic range to see oscillations of these pilot bunches when configured for nominal intensities. Already at this stage a large amount of data can therefore be automatically removed.

The second processing step is to align the data points corresponding to a given bunch over all of the recorded turns. As the sampling rate of the digitizers is not a multiple of the beam's revolution frequency, there is a noninteger number of samples between the bunch in two subsequent turns. In addition, although the particles in the LHC are already highly relativistic, the revolution frequency nevertheless changes with energy. For the LHC the difference in frequency between injection and collision energy is 0.024 Hz, corresponding to a difference of 1.926 samples/turn for the 10 GSPS acquisition system. As it is not possible to know the LHC revolution frequency at the precise moment of acquisition, and since the free-running sampling clock in the digitizer may drift over time, a fixed correction cannot be used and the number of samples per turn must be determined from the data. This is done by comparing the sum signal of the first populated bunch slot and using an iterative algorithm to find the best fit between two turns. First, an approximate value for the revolution frequency is used, with the overlap between two consecutive turns varied in coarse steps to search for the fit which minimizes the average squared difference between the corresponding samples on each turn. Then, to increase the precision, the scan is repeated for smaller step sizes around the best fit value initially obtained.

Once the data is aligned, the last processing step is to remove a large "baseline" on the difference signal which is constant from turn to turn. This signal is a result of the combination of a number of effects: (1) Beam orbit offset from the pickup center, resulting in a difference signal proportional to the beam position. (2) Differences in cable attenuation and delay before the hybrid, leading to an artificial orbit offset. (3) Imperfections in the hybrid, resulting in "leakage" from the sum signal into the difference signal even for a well-centered beam [26].

While the first two effects can be minimized by careful control of the orbit and optimization of the cabling, the third is impossible to avoid but can be removed during postprocessing. For instability measurements we are interested in the oscillation signal which varies from turn to turn but, over many turns, is expected to average out. Therefore



FIG. 6. A mode 4 instability measured in the LHC showing the raw difference signal before (top) and after (bottom) the baseline correction has been applied. For reference, the dotted line shows the sum signal of the same bunch.

the baseline, which is constant from turn to turn, can be obtained by calculating the sample by sample mean of all acquired turns. This mean can then be subtracted from each turn to remove the baseline. For example, for a sample a_n in turn *t* the corrected sample $a_{n,\text{corr}}$ is obtained by

$$a_{n,\text{corr}}(t) = a_n(t) - \sum_{\tau=1}^T \frac{a_n(\tau)}{T},$$

where *T* is the number of turns of data acquired. This procedure is complicated by the fact that the number of samples per turn is not an integer, since the sampling frequency is not correlated with the revolution frequency. Before the average can be made, the data from each turn needs to be resampled to common sampling points using a linear interpolation. The calculated mean values are then interpolated back to the original sampling points of each turn before subtraction. An example of a mode |m| = 4 instability, both before and after the baseline removal, is shown in Fig. 6.

Although the baseline removal works well, the uncorrected amplitude of the raw signal is still digitized and reduces the dynamic range available for the oscillation signal. This can be seen in Fig. 6 where the uncorrected signal spans almost the entire dynamic range of the digitizer with the oscillation signal occupying less than half of the full scale range. Improved hybrids, in order to reduce this residual signal, are an upgrade that could lead to significant improvements to the overall system sensitivity. However, given the bandwidth and power-handling requirements, no suitable alternative has been identified to date. Novel methods using electro-optical detection [27] and frequency



FIG. 7. Calibration of the head-tail monitor response using orbit bumps in the pickup in both planes of beam 2.

domain analysis [28,29] are currently being studied as potential future techniques to overcome this limitation.

In order to give an oscillation value in units of microns, the head-tail monitors have been calibrated by performing orbit bumps in the pickups. The values of the sum and difference signals are taken at the sample corresponding to the center of the bunch and the normalized position is calculated as the difference divided by sum, to give a value that is proportional to the position but independent of beam intensity. The results from such a calibration are plotted in Fig. 7 against the programmed bump position for a bump range of ± 5 mm around the reference orbit. From this data, a linear fit finds that a normalized position of 0.042 corresponds to 1 mm. The small offset observed, which is different in each plane, comes from the fact that the reference orbit is not completely centered in the pickup. Since the head-tail monitor only needs to measure relative position changes this offset can be ignored, although it does contribute to the baseline residual.

During regular operation the head-tail monitor is triggered automatically via the LIST network. As the instrument that generates the trigger may have a much higher sensitivity than the head-tail monitor, this can lead to triggers where the oscillations are at a level that cannot be seen. In order to automatically discard such datasets it is necessary to determine if an acquisition contains a visible instability. A simple approximation of the amplitude of the instability can be determined by looking at the ratio of the amplitude of the baseline corrected difference signal inside and outside the bunch area, with the bunch area determined by locating the points of the sum signal above a certain amplitude threshold. If this ratio is higher than a certain threshold value, the bunch is considered to contain an instability. In addition to the amplitude of the instability, the mode number can also be determined by searching for the zero-crossing points of the envelope of the difference signal within the bunch.

The information on the oscillation amplitude and instability mode number can be used to automatically categorize acquisitions obtained during regular operation. Those which contain "interesting" events, for example large oscillation



FIG. 8. The operational principal of the LHC "base-band tune" (BBQ) system.

amplitudes, can be identified and studied while others can be automatically removed to save disk space.

IV. LHC TUNE MEASUREMENT SYSTEM

The betatron tune is measured in the LHC using the baseband tune (BBQ) system [30]. The analog front end, shown in Fig. 8, is based around diode peak detectors directly connected to the electrodes of a transverse BPM. The betatron oscillations of the beam result in an amplitude modulation of each BPM electrode signal at the tune frequency q and the peak detectors directly extract this low frequency, small amplitude, modulation signal from the high-frequency, large amplitude, bunch signal. The large common mode signal is removed using dc blocking capacitors and the difference of each pair of electrodes is performed in base-band using a precision operational amplifier. A sharp notch filter suppresses the strong 11.245 kHz revolution frequency component, allowing the betatron frequency to be amplified without saturating the electronics. In the LHC, where $q \approx 0.3$, the frequency range of interest extends from approximately 1 to 5 kHz and, as the particles are already highly relativistic, does not change much during acceleration. The tune signal can therefore be sampled with a high-resolution 24 bit audio ADC that is frequency locked to a multiple of the revolution frequency using a phase locked loop. A field programmable gate array (FPGA) based digital acquisition system performs a fast Fourier transformation (FFT) of the ADC data [31,32], with the tune value extracted from the resulting spectrum using a peak-detection algorithm with spectral interpolation to improve the resolution [33].

Because of its extreme sensitivity to oscillations (down to tens of nanometers), the BBQ system is likely to detect the onset of an instability well before any other instrument and can therefore serve as a trigger source for less sensitive instruments such as the head-tail monitor. An example of a typical single bunch instability recorded by the BBQ in the LHC is shown in Fig. 9(a). The increase in the amplitude of the spectral component corresponding to the tune frequency is evident during the instability. Figure 9(b) shows an acquisition from the head-tail monitor well after the BBQ starts to show enhanced activity, revealing intrabunch motion with approximately 200 μ m oscillation amplitude.

While tracking the amplitude of the tune peak provides a good indication of the growth of an instability, in order to provide an online trigger the processing time must be minimized and the overhead of performing the FFT and peak detection is significant. To avoid this latency, and simplify the processing, the amplitude growth of the time domain signal can be used as a substitute, with the assumption that the tune is the dominant spectral component in the BBQ signal.

The first version of an instability trigger algorithm for the BBQ system, called the "three-averages" algorithm, was developed in 2013 based on simulated and recorded beam data [34]. For the start of run 2 in 2015, this trigger algorithm was implemented in an FPGA connected to the BBQ acquisition system. Having the algorithm running in real time helped evaluate its performance under operational beam conditions. Although the algorithm was found to perform well under specific conditions, during standard physics operation it proved to be extremely sensitive to small amplitude fluctuations, which resulted in a large number of spurious triggers. A second algorithm, the "record-values" algorithm, has since been developed using a different principle which attempts to reduce the number of false triggers under operational conditions.

A. Algorithms for detecting instabilities

1. Three-averages algorithm

In order to detect that an instability is occurring, it is necessary to look for a growth in the amplitude of the betatron oscillation of the beam. Defining an absolute amplitude threshold, above which the beam is considered unstable, is not sufficient as the oscillation may have a stable, constant amplitude. It therefore also needs to take



(a) LHC tune system time domain data (top), frequency spectrogram (middle) and tune peak amplitude (bottom) for a typical instability that occurred between 02:45 and 02:55 on 21st May 2017. The dotted line shows the time of the Head-Tail measurement that is shown in Fig. 9b.



(b) Superposition of 11 turns of the Head-Tail difference signal at 02:50:38 showing the instability mode.

FIG. 9. A typical single bunch instability seen by the LHC tune measurement system in both time and frequency domain, and an acquisition of the head-tail monitor during the same event.

into account that an instability implies a rising amplitude. As a measure of the oscillation amplitude, the threeaverages trigger algorithm computes a moving average of the standard deviation (σ) about the mean of the BBQ signal for three different time windows. The window lengths are dimensioned such that

$$W_{\rm short} < W_{\rm med} < W_{\rm long}$$
.

Under stable beam conditions, it is expected that the amplitude of all windows will be approximately equal, so

$$\bar{\sigma}_{\rm short} \approx \bar{\sigma}_{\rm med} \approx \bar{\sigma}_{\rm long}$$

In contrast, during an instability, the amplitude of the input signal increases and the moving averages of the standard



FIG. 10. BBQ time domain data and moving averages illustrating their behavior during a typical instability.

deviation change with a rate that corresponds to the length of their windows, as shown in Fig. 10. In this case, the following inequalities will hold:

$$\begin{split} \bar{\sigma}_{\rm short} &- \alpha \bar{\sigma}_{\rm med} > 0 \\ \bar{\sigma}_{\rm med} &- \beta \bar{\sigma}_{\rm long} > 0, \end{split}$$

where $\alpha, \beta > 1$ are coefficients chosen to reduce the influence of noise. Fulfilling these conditions for many consecutive turns is a clear indicator of the growth of an instability.

To detect the instability, a counter (C_i) is incremented on each turn by the normalized difference of each pair of window functions:

$$C_i = C_{i-1} + w_{\alpha} \frac{\bar{\sigma}_{\text{short}} - \alpha \bar{\sigma}_{\text{med}}}{\bar{\sigma}_{\text{short}} + \bar{\sigma}_{\text{med}}} + w_{\beta} \frac{\bar{\sigma}_{\text{med}} - \beta \bar{\sigma}_{\text{long}}}{\bar{\sigma}_{\text{med}} + \bar{\sigma}_{\text{long}}},$$

where w_{α} , w_{β} are weighting factors corresponding to the number of turns required to confirm the presence of an instability for this window. Once the counter reaches a threshold value, a trigger is generated and the counter is reset to zero. In order not to generate a large number of consecutive triggers, the counter is held at zero for a short "hold-off" period after each trigger.

2. Record-values algorithm

The stable oscillation of the BBQ signal can be viewed as a sequence (Q) of *n* values oscillating independently around a constant value. From this, a subsequence (S)can be defined that consists of only the historical maximum values, i.e., each new maximal value encountered when iterating through the sequence Q. It can be shown [35] that if the length of S is treated as a random variable, for large values of *n* the expected value (μ) and standard deviation (σ) of this random variable are as follows:

$$\mu = \sum_{i=1}^{n} \frac{1}{i},$$

$$\sigma = \sqrt{\sum_{i=1}^{n} \frac{1}{i} - \sum_{i=1}^{n} \frac{1}{i^{2}}}$$

which, for large values of n, may be approximated by

$$\begin{split} \mu &\approx \ln(n) + \gamma, \\ \sigma &\approx \sqrt{\ln(n) + \gamma - \frac{\pi^2}{6}} \end{split}$$

where γ is the Euler-Mascheroni constant:

$$\gamma = \lim_{n \to \infty} \left(-\ln(n) + \sum_{k=1}^{n} \frac{1}{k} \right) \approx 0.5772$$

Compared to this steady state result, the situation changes dramatically if there is a growth in the amplitude of the signal. In this case, it is expected that almost every local peak value in the sequence will be a new historical maximum "record value." Therefore, a simple method to determine whether the beam becomes unstable is to count the number of record values in the last k samples and compare it to the theoretical expectation.

A slight modification of the approach enables the algorithm to run in amortized constant time, meaning that the average time required to process each new sample does not depend on k nor on any other parameter of the algorithm. The modification is to count the number of record values appearing in the sequence until there has not been a new one in the last d samples. If the number of record values reaches a level which is unlikely from the theoretical predictions then a trigger is generated. On the contrary, if a new record value is not encountered after d samples, then the counter is reset and the highest value in the last d samples becomes the starting record value for the next sequence.

To explain of the operation of the record-values algorithm, a simplified example is presented in Fig. 11 where the simulated signal represents the peak of the BBQ signal during an instability event. During the first 20 turns the simulated signal is stable. New local maximum values are encountered at turns 3 and 13 and result in the record-value counter being incremented on these turns. After turn 13, a new local maximum value is not found in the next d = 4turns and, at turn 18, the maximum value in these turns is taken as the new local maximum and the counter is reset to zero. Starting at turn 20, there is an exponential amplitude growth introduced into the simulated signal. Accordingly, almost every turn is a new local maxima and the counter increments quickly. Once the threshold on the counter value is passed, a trigger is generated on each new turn with a larger amplitude, as indicated by the markers. After turn 40, the simulated signal returns to a stable state. Both the local maximum and the counter remain at their maximal value for d = 4 turns. After this period, the counter is again reset to zero and the highest amplitude of the signal during turns 40-44 becomes the new local maximum and the process repeats.



FIG. 11. Simulation of the peak of the BBQ signal during an instability to illustrate the operation of the record-values algorithm with d = 4 and a trigger threshold of 4.

In order to reduce the probability of a trigger caused by the large transients that can occur during beam injection, a further check is made on each sample. If the new peak value is significantly greater than the maximum in the last ksamples it is indicative of a transient rather than an instability and is therefore ignored.

B. Comparison between the algorithms

In order to compare their performance, both algorithms were applied to data recorded during known instabilities under various beam conditions.

One important improvement that was implemented in the record-values algorithm is to avoid triggering on injection transients. The behavior of the two algorithms to an injection transient is shown in Fig. 12. The three-averages algorithm triggers on the large increase of amplitude caused by the injection event as $\bar{\sigma}_{short}$ increases rapidly above $\bar{\sigma}_{med}$ and $\bar{\sigma}_{long}$ with the counter quickly reaching its threshold. Due to the hold-off period after each trigger, the three-averages algorithm is not able to react to the following instability. In comparison, the record-values algorithm



FIG. 12. Comparison of the trigger algorithms during the transient caused by an injection.



(a) BBQ time-domain data recorded during the instability



(b) Head-Tail Monitor acquisition at the peak of the instability

FIG. 13. Comparison of the trigger algorithms during an instability with a fast rise time that occurred with few bunches in the machine on the 31st of July in 2016.

filters the transient and is able to trigger on a subsequent rise in BBQ amplitude which could be indicative of an instability.

In cases where there are only a few bunches in the machine and the transverse feedback gain is low, the signals from the BBQ are clean with a sharp tune peak, as shown in Fig. 15(a). Under these conditions any bunch becoming unstable has a large effect on the amplitude of both the time and frequency domain data. When the instability has a reasonably fast rise time, both algorithms perform in a similar way, as seen in Fig. 13. It is notable that for this measurement, although the BBQ amplitude growth is clear, the corresponding head-tail monitor measurement, shown in Fig. 13(b), is not particularly good. While it is possible to see some increase in amplitude in the region of ± 0.5 ns, corresponding to the bunch area, it is hard to determine the mode of the oscillation.

An example with a slower rise time is shown in Fig. 14. In this case, the three-averages algorithm does not trigger as the averages of all three windows increase at a similar rate so that there is never a large enough difference between them to cause a trigger. As the record-value algorithm only considers a growth in amplitude, and is not affected by its rate, the instability is detected and generates triggers during the event with the time between triggers given by a userdefined hold-off period.

During normal physics fills, the LHC operates with a large number of bunches and relies on a high transverse feedback gain for stability. In these conditions the tune peak is broadened and the BBQ spectrum can become dominated by other frequency components, as shown in Fig. 15(b).



FIG. 14. Comparison of the trigger algorithms during an instability with a slow rise time and few bunches in the machine that occurred on the 31st of July in 2016.

The large frequency component at around 0.1 is thought to be related to longitudinal beam motion that adds additional side bands at the revolution frequency and which are mixed down to dc. These additional noise components make the determination of the tune extremely challenging. To mitigate this a second BBQ system is used for operational tune measurement system, which is gated to measure only on a few "witness" bunches that have lower transverse feedback gain. However, as the instability trigger must be able to detect an instability on any bunch in the machine, it cannot use the gated system and must be able to cope with the additional noise components.

The degradation of the tune signal quality has clear consequences for the performance of the trigger algorithms as noise and other artifacts make interpretation of the signal more challenging. In addition, a single bunch, or a small



FIG. 15. Comparison of typical BBQ spectra observed under different beam conditions. With a small number of bunches and low transverse feedback gain (top) the tune peak at 0.295 is clearly visible. With a full machine and high transverse feedback gain (bottom) the peak is less visible and other noise sources become dominant.



FIG. 16. Instabilities with 590 nominal bunches recorded on the 1st of August in 2016 that are seen by the three-averages algorithm but not by the record-values algorithm.

number of bunches, becoming unstable out of several hundred has a much smaller effect on the overall amplitude of the time domain signal. While the record-value algorithm handles the additional noise quite well, it cannot detect the small amplitude variations caused by single bunches becoming unstable. The more sensitive three-averages algorithm, on the other hand, causes a large number of spurious triggers but is able to detect the small changes in amplitude linked to instabilities. One such example is shown in Fig. 16. In this case, the head-tail monitor acquisition confirmed that only two out of the 590 circulating bunches had become unstable. Although the oscillation amplitude of these single bunches was almost 300 μ m, their overall contribution to the BBQ tune signal amplitude can be seen to be very small.

During the course of LHC run 2, both of the trigger algorithms presented have been used operationally under a wide variety of beam conditions. While both algorithms perform well, and each has advantages under certain beam conditions, neither one suits all operational scenarios. In addition, both suffer from the poor tune signal quality that is typically observed during physics fills which limits any further improvement to their performance. Despite these limitations, both have been used successfully to identify the instabilities that have occurred during regular operation throughout run 2. For LHC run 3, a parallel implementation of the two algorithms is planned to allow the optimal algorithm to be selected depending on the beam conditions.

The algorithms presented have been designed to run in real time on a FPGA with minimal trigger latency and, because of this, are limited to operating on time-domain data. More complex algorithms could be envisaged, for example ones that look for instabilities in the frequency domain. By looking only at the tune frequency, such algorithms have the potential to be more robust against interference from the additional noise components seen in the tune spectrum. Due to the processing overhead of frequency-domain analysis, it would necessarily have a higher latency than the time-domain algorithms. However, as many instability events observed in the LHC have quite slow rise times, this may be acceptable if the algorithm could be made more robust. The extensible nature of the LHC instability trigger network also allows additional trigger sources to be added, so the possibility of using other instruments to detect the growth of instabilities is being investigated as another potential future extension of the instability trigger system.

V. CONCLUSION

This paper presents the newly developed transverse beam instability detection and monitoring systems for the LHC. After several years of operation, the systems are now showing excellent performance and have become essential tools to identify, combat and cure transverse beam instabilities that occur during regular machine operation.

The detection system is based on the observation of the growth of submicron transverse beam oscillations using the base-band tune system. An online detection algorithm sends out a trigger via the dedicated LHC instability trigger network allowing other systems to capture data relevant to the event. Two different trigger algorithms have been developed and tested in order to identify the onset of an instability. The performance of these algorithms is highly dependent on the tune signal quality and hence they have been tuned to minimize the number of false triggers due to noise, while still being sensitive enough to be able to identify small instability events. Although neither algorithm is ideal for all beam conditions, they have both been used to successfully identify instabilities that have occurred during LHC run 2.

One instrument that profits from this triggering network is the LHC head-tail monitor, which is able to acquire transverse oscillations within individual LHC bunches. The head-tail monitor has been shown to be able to detect instabilities up to mode 4 with typical beam conditions in the LHC. Recent upgrades enable it to acquire up to 450 turns for all 2808 bunches with a detectable resolution of less than 100 μ m. Each event represents a large quantity of data that must be postprocessed in order to identify which bunches in the machine are becoming unstable as well as the instability mode. In order to reduce the amount of data which must be stored and manually evaluated, automatic postprocessing of the data has been implemented to store only the acquisitions containing instabilities. Future upgrades are now under investigation to design a higher bandwidth pickup and provide higher sensitivity detection systems.

- C. Pellegrini, A. Marinelli, and S. Reiche, The physics of x-ray free-electron lasers, Rev. Mod. Phys. 88, 015006 (2016).
- [2] C. Bostedt, S. Boutet, D. M. Fritz, Z. Huang, H. Ja Lee, H. T. Lemke, A. Robert, W. F. Schlotter, J. J. Turner, and G. J. Williams, Linac coherent light source: The first five years, Rev. Mod. Phys. 88, 015007 (2016).
- [3] T. Wyatt, High-energy colliders and the rise of the standard model, Nature (London) **448**, 274 (2007).
- [4] A. W. Chao, *Physics of Collective Beam Instabilities in High Energy Accelerators* (Wiley, New York, 1993).
- [5] S. A. Heifets and S. A. Kheifets, Coupling impedance in modern accelerators, Rev. Mod. Phys. 63, 631 (1991).
- [6] K. Ohmi and F. Zimmermann, Head-Tail Instability Caused by Electron Clouds in Positron Storage Rings, Phys. Rev. Lett. 85, 3821 (2000).
- [7] R. Cimino, I. R. Collins, M. A. Furman, M. Pivi, F. Ruggiero, G. Rumolo, and F. Zimmermann, Can Low-Energy Electrons Affect High-Energy Physics Accelerators?, Phys. Rev. Lett. 93, 014801 (2004).
- [8] E. Benedetto, G. Franchetti, and F. Zimmermann, Incoherent Effects of Electron Clouds in Proton Storage Rings, Phys. Rev. Lett. 97, 034801 (2006).
- [9] G. Rumolo, G. Arduini, E. Métral, E. Shaposhnikova, E. Benedetto, R. Calaga, G. Papotti, and B. Salvant, Dependence of the Electron-Cloud Instability on the Beam Energy, Phys. Rev. Lett. **100**, 144801 (2008).
- [10] E. Métral *et al.*, Beam instabilities in hadron synchrotrons, IEEE Trans. Nucl. Sci. **63**, 1001 (2016).
- [11] J. Gareyte and F. Sacherer, Head-tail type instabilities in the CERN PS and Booster, in *Proceedings of the 9th International Conference on High-Energy Accelerators* (*HEACC'74*), *Stanford, California, 1974*, pp. 341–346.
- [12] A. Hofmann, *The Physics of Synchrotron Radiation*, Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology (Cambridge University Press, Cambridge, England, 2004).
- [13] E. Rossa, Real time single shot three-dimensional measurement of picosecond photon bunches, AIP Conf. Proc. 333, 148 (1995).
- [14] E. Métral *et al.*, Measurement and interpretation of transverse beam instabilities in the CERN Large Hadron Collider (LHC) and extrapolations to HL-LHC, in *Proceedings of the 57th ICFA Advanced Beam Dynamics* Workshop High-Intensity and High-Brightness Hadron Beams (HB'16), Malmö, Sweden, pp. 254–259.
- [15] E. Métral *et al.*, Review of the instabilities observed during the 2012 run and actions taken, in *Proceedings of the 4th Evian Workshop LHC Beam Operation, Evian-les-Bains, France, 2012*, pp. 87–94.
- [16] G. Iadarola, H. Bartosik, K. Li, L. Mether, A. Romano, G. Rumolo, and M. Schenk, Performance limitations from electron cloud in 2015, in *Proceedings of the 6th Evian*

Workshop LHC Beam Operation, Evian-les-Bains, France, 2015, pp. 101–110.

- [17] A. Mereghetti *et al.*, β^* -Reach IR7 Collimation Hierarchy Limit and Impedance, CERN ACC Note 2016-0007, 2016 (unpublished).
- [18] A. Lechner *et al.*, BLM thresholds for post-LS1 LHC operation: UFOs and orbit bumps in the arcs and straight sections, in Proceedings of the Workshop Beam Induced Quenches, CERN, Geneva, Switzerland, 2014 (unpublished).
- [19] L. R. Carver *et al.*, Current status of instability threshold measurements in the LHC at 6.5 TeV, in *Proceedings of the 7th International Particle Accelerator Conference* (*IPAC'16*), Busan, Korea, 2016, pp. 1434–1437.
- [20] D. Domínguez, J. J. Gras, J. Lewis, J. J. Savioz, J. Serrano, and F. J. Ballester, An FPGA based multiprocessing CPU for beam synchronous timing in CERN's SPS and LHC, in *Proceedings of the 9th International Conference Accelerator and Large Experimental Physics Control Systems* (ICALEPCS'03), Gyeongju, Korea, 2003, pp. 113–115.
- [21] J. Serrano et al., The White Rabbit project, in Proceedings of the 12th International Conference Accelerator and Large Experimental Physics Control Systems (ICA-LEPCS'09), Kobe, Japan, 2009, pp. 93–95.
- [22] T. Włostowski, J. Serrano, G. Daniluk, M. Lipiński, and F. Vaga, Trigger and RF distribution using White Rabbit, in Proceedings of the 15th International Conference Accelerator and Large Experimental Physics Control Systems (ICALEPCS'15), Melbourne, Australia, 2015, pp. 619–623.
- [23] D. Cocq, O. R. Jones, and H. Schmickler, The measurement of chromaticity via a head-tail phase shift, AIP Conf. Proc. 451, 281 (1998).
- [24] S. D. Fartoukh and R. Jones, Determination of chromaticity by the measurement of head-tail phase shifts: Simulations, results from the SPS and a robustness study for the LHC, CERN LHC Project Report 602, 2002 (unpublished).
- [25] R. E. Schafer, Beam position monitoring, AIP Conf. Proc. 249, 601 (1992).
- [26] G. Kotzian, W. Hofle, R. J. Steinhagen, D. Valuch, and U. Wehrle, Evaluation of strip-line pick-up system for the SPS wideband transverse feedback system, in *Proceedings* of the 2nd International Beam Instrumentation Conference (IBIC'13), Oxford, UK, 2013, pp. 690–693.
- [27] S. M. Gibson, A. Arteche, S. E. Bashforth, A. Bosco, M. Krupa, and T. Lefevre, Enhanced bunch monitoring by interferometric electro-optic methods, in *Proceedings of the 9th International Particle Accelerator Conference (IPAC'18), Vancouver, Canada, 2018*, pp. 2353–2356.
- [28] R. J. Steinhagen, M. J. Boland, and T. G. Lucas, A multiband-instability-monitor for high-frequency intra-bunch beam diagnostics, in *Proceedings of the 2nd International Beam Instrumentation Conference (IBIC'13), Oxford, UK*, 2013, pp. 327–330.
- [29] T. E. Levens, T. Lefevre, and D. Valuch, Initial results from the LHC multi-band instability monitor, in *Proceedings of the 7th International Beam Instrumentation Conference* (*IBIC'18*), *Shanghai, China, 2018*, pp. 314–318.
- [30] M. Gasior, Faraday cup award: High sensitivity tune measurement using direct diode detection, in *Proceedings*

of the 2012 Beam Instrumentation Workshop (BIW'12), Newport News, Virginia, USA, 2012, pp. 1–8.

- [31] A. Boccardi, M. Gasior, O. R. Jones, K. K. Kasinski, and R. J. Steinhagen, The FPGA-based continuous FFT tune measurement system for the LHC and its test at the CERN SPS, in *Proceedings of the 22nd Particle Accelerator Conference (PAC'07), Albuquerque, New Mexico, USA*, 2007, pp. 4204–4206.
- [32] M. Gasior and A. Boccardi, High dynamic range spectral analysis in the kHz band, in *Proceedings of the 9th European Workshop Beam Diagnostics and Instrumentation for*

Particle Accelerators (DIPAC'09), Basel, Switzerland, 2009, pp. 348–350.

- [33] M. Gasior, Improving frequency resolution of discrete spectra, Ph.D. thesis, AGH University of Science and Technology, 2006.
- [34] J. Ellis, Level-1 trigger development for the multi-band instability monitor, CERN Students Note 2013-103, 2013 (unpublished).
- [35] N. Glick, Breaking records and breaking boards, Am. Math. Mon. 85, 2 (1978).