INSTABILITY DIAGNOSTICS

T. Levens, A. Boccardi, A. Butterworth, L.R. Carver, G. Daniluk, M. Gasior,
W. Höfle, O.R. Jones, G. Kotzian, T. Lefevre, J.C. Molendijk, M. Ojeda Sandonís,
J. Serrano, D. Valuch, T. Włostowski, CERN, Geneva, Switzerland
F. Vaga, University of Pavia, Italy

Abstract

This paper gives an overview of the available systems for transverse instability diagnostics in the LHC with a focus on those that have been added or upgraded since run 1. The status and performance during 2015 will be reviewed. Finally the outlook for 2016 and planned improvements will be discussed.

INTRODUCTION

Run 1 of the LHC saw transverse instabilities at injection and during physics fills with 50 ns bunch spacing [1]. The change to 25 ns bunch spacing, combined with stricter limits on beam loss at 6.5 TeV, mean that the mitigation of instabilities will also be an important consideration during run 2. Understanding and characterising instabilities is fundamental to ensure that the correct adjustments can be made to the machine settings.

An internal review on instability diagnostics [2] was held on 15th March 2013 and identified a number of areas that could be improved for run 2. During the recent long shutdown (LS1) there have been a number of developments and improvements to the instability diagnostics available for the LHC, which will be discussed in this paper. Their performance during the 2015 run will be evaluated and an outlook will be provided for 2016.

LHC INSTABILITY TRIGGER NETWORK

The LHC Instability Trigger Network (LIST) has been deployed during LS1 and made operational during 2015. The network allows bi-directional trigger distribution between the various instruments capable of detecting transverse instabilities such that, if one detects an instability, it can synchronously trigger other measurement devices.

The LHC is served by two distinct timing networks, the millisecond synchronous "General Machine Timing" (GMT)
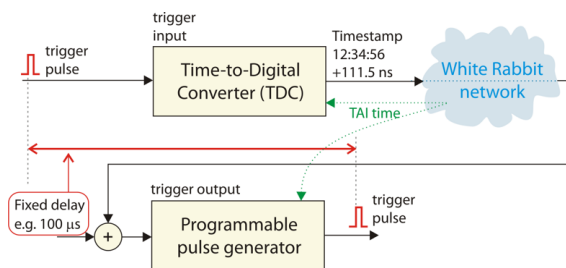

Figure 1: A diagram showing the principal of trigger distribution over a White Rabbit network.


Figure 2: LIST network deployed during LS1.

used by all systems and turn synchronous "Beam Synchronous Timing" (BST) dedicated to beam instrumentation. The option to distribute triggers with these existing networks was considered. However, both networks are unidirectional and are designed to distribute centrally generated timing events to equipment. Therefore, they offer limited possibilities to receive and redistribute triggers generated by remote equipment. Furthermore, both networks have limited bandwidth available to distribute events. Because of these limitations, a new network has been deployed that is dedicated to distribution of instability trigger events.

The LIST network is based on White Rabbit [3] technology, a deterministic, synchronous extension to the Ethernet standard. White Rabbit allows geographically separated nodes on a network to be synchronised with sub-nanosecond precision. The basic operation of the LIST network is shown in Fig. 1 and is described in more detail in [4]. A trigger pulse, generated by an instrument, has its arrival time tagged by a time-to-digital converter (TDC). The generated timestamp is distributed over the network and a pulse can be generated at any number of outputs after a fixed delay. As all nodes on the network share a synchronised clock, and hence a common notion of time, the input to output delay can be precisely programmed down to a minimal value imposed by the propagation delay of the network. For the LHC, the minimum delay is around 300 μs due to the long fiber optic links between the CCR and Point 4.

The LIST network is based on BE-CO's standard FMC hardware kit [5] and was fully deployed during LS1. As shown in Fig. 2, nodes have been installed in SR4 & UX45 for BE-RF and UA43 & UA47 for BE-BI with a master node in the CCR for connection to the central timing systems.

During 2015 the first instruments were connected and commissioned, starting with inputs from the tune measurement system (BBQ) and outputs to the Head-Tail monitor. Later in the year, outputs were connected to the Transverse Damper (ADT). Connections have also been made in order to redistribute events through the GMT and BST networks, enabling existing instrumentation to be triggered synchronously.

The LIST network has performed well during 2015 and is now routinely used to trigger the Head-Tail monitor from the BBQ. All trigger events are stored in the logging database to help with diagnosis of problems with the network and to facilitate the tuning of the trigger algorithms. For the 2016 start-up it is planned to fully commission the connection to the ADT Observation Box. As the network has a fully modular design in both hardware and software, it will possible to connect additional instruments as required during the run.

DEVELOPMENT OF BI INSTRUMENTS

Head-Tail Monitor

The LHC Head-Tail monitor is based on the fast acquisition of a stripline beam-position monitor (BPM). A sum (Σ) and difference (Δ) signal for each plane is calculated using a wideband RF hybrid and directly digitised by a 10 GSPS, 8 bit digitiser. Although originally foreseen for chromaticity measurements [6], due to the digitizer's high sampling rate, the system can be used for direct time-domain measurements of intra-bunch motion. However, the minimum detectable oscillation amplitude is limited by the digitizer's dynamic range and, due to the high sampling rate, the maximum acquisition length is limited to 11 turns. These two factors require the Head-Tail monitor to be precisely triggered once the oscillation amplitude has reached a sufficient level to be visible. The large amount of data produced per acquisition also poses challenges for storage and processing of the data.

Since run 1 there have been no major hardware changes, although the digitizers have been relocated to a new rack in UA47 and the pick-up cabling has been checked and optimised. The acquisition software has been completely rewritten with a new FESA 3 class and new GUIs that have been tailored for the system's operation as an instability monitor. Acquisitions are saved directly to disk for offline analysis and the operational GUI provides data post-processing to allow identification of instability modes.

A major limitation during run 1 had been that the Head-Tail monitor could only acquire one plane (horizontal or vertical) at a time. The cabling optimisation and new FESA class have removed this limitation, allowing both planes to be acquired simultaneously.

The Head-Tail monitor has been used extensively during the 2015 MD blocks and has performed reliably with no major issues. Some selected results, showing examples of modes $|m| = 0$ to $|m| = 4$, are shown in Fig. 3.

BBQ Trigger

A new development during 2015 has been the addition of an instability trigger to the LHC tune measurement system

(a) Mode $|m| = 0$

(b) Mode $|m| = 1$

(c) Mode $|m| = 2$

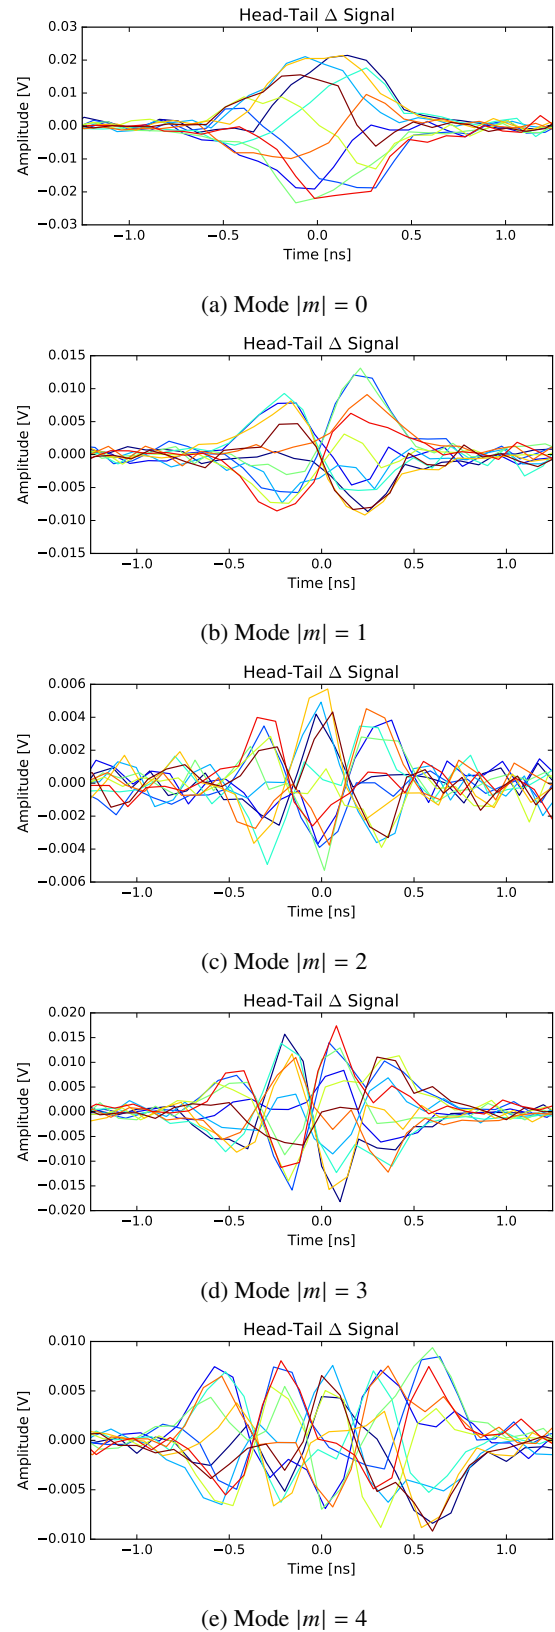(d) Mode $|m| = 3$

(e) Mode $|m| = 4$

Figure 3: Examples of instabilities measured with the LHC Head-Tail monitor during MD sessions in 2015.

Figure 4: Example of an instability from the BBQ Trigger system. The top axis shows the raw time domain BBQ signal and the bottom axis shows the generated triggers with red markers.

Figure 5: Measurement of MIM filter bank bands (solid lines) compared to simulated values (dotted lines).

(BBQ). As the BBQ is an extremely sensitive detector, it has the chance to detect the onset of an instability before any other instrument. Therefore, it can serve as a trigger source for less sensitive instruments connected to the LIST network. In order to be low-latency, the trigger algorithm [7] looks for a growth in the amplitude of the time-domain BBQ data, indicative of an instability. The algorithm runs in real time on a dedicated FPGA and is installed in parallel with the standard BBQ acquisition system. During 2015, the trigger was installed on the development BBQ system on the surface in SX4. For the 2016 run it is planned that it will be moved to the operational, high-sensitivity, BBQ systems in UA47 & UA43.

Although the initial implementation of the algorithm correctly generates triggers on instabilities, as shown in Fig. 4, it has been found to be extremely sensitive to noise and small perturbations of the BBQ signal. It also generates false triggers on injection and dump transients and during ADT excitation for abort gap and injection cleaning. As each trigger can potentially generate a lot of data from other monitors, improvements to the algorithm will be studied for the 2016 run in order to improve its robustness and to mask transients and excitation.

Multi-Band Instability Monitor

The Multi-Band Instability Monitor (MIM) is an ongoing hardware development in BE-BI as a detector for intra-bunch instabilities. Instead of direct time-domain digitisation, as used by the Head-Tail monitor, the BPM signal is processed in the frequency domain. An RF band-pass filter bank is used to select different frequency components of the signal which can then be sampled independently with higher resolution analogue to digital converters. By looking at the relative power in each band, it is possible to determine the instability mode number [8].

The first prototype of the MIM hardware has been installed in the LHC at the end of 2015. Its filter bank, developed in

collaboration with BE-RF, features eight frequency bands at multiples of the LHC RF (400 MHz to 3.2 GHz). A measurement of the first prototype filter bank is shown in Fig. 5 which, for the bands above 1.2 GHz, shows good agreement with simulation and excellent isolation between the bands. The degraded response of the 800 MHz band is still under study, but is not expected to be a problem for the initial tests. Detection of each band is performed using diode detectors sampled with high-resolution ADCs. The architecture, similar to the BBQ provides a high sensitivity but does not allow bunch-by-bunch measurements.

Currently a single MIM front-end has been installed and is switchable in order to measure a single plane of interest. Some preliminary tests have been made during the 2015 ion run and the first measurements with protons are expected early in 2016.

ADT OBSERVATION BOX

The LHC Transverse Damper (ADT) can provide high-resolution bunch-by-bunch normalised beam position information from each of its pick-ups. During run 1, the amount of data which could be captured was limited to 73 turns (for all bunches) by the size of the SRAM memories available on the ADT VME processing board.

As part of the substantial upgrade to the ADT, carried out during LS1 [9], a general purpose "Observation Box" has been developed by BE-RF [10]. This system uses a commercial server PC and the BE-CO developed "SPEC" [11] card with custom firmware developed by BE-RF. The SPEC card captures the 40 MSPS bunch-by-bunch beam position data that is sent across gigabit serial links between the ADT Beam Position Module and Digital Signal Processing Unit as shown in Fig. 6. Inside the server, the data is stored into a large memory buffer where it can be made available for analysis.

An initial deployment of the Observation Box was made for the longitudinal RF system to capture the raw I/Q data from the cavity sum and phase pick-up. From these signals, the calculated bunch-by-bunch stable phase has been used
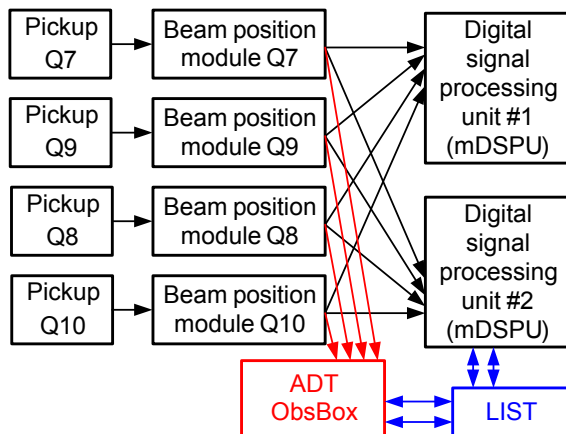
Figure 6: ADT block diagram after LS1. ADT internal signals are marked in black, connections to the Observation Box in red and to the LIST in blue.

Figure 7: Injection oscillations for a 72 bunch train captured by the ADT Observation Box.

extensively for electron-cloud diagnostics during the 2015 scrubbing runs [12].

Towards the end of the 2015 run, the first connections from the ADT to the Observation Box were made in order to test the FESA class and allow development of analysis scripts. As an Observation Box can process four signals, initially only the Q7 horizontal and vertical pick-ups for both beams were connected to a single server [13]. For the 2016 start-up it is planned to deploy three additional servers to connect the existing Q9 pick-ups and the newly introduced Q8 and Q10 pick-ups.

An initial version of the FESA class for the Observation Box has been deployed during 2015. The class implements a single rolling buffer of 4.2 million turns, equivalent to 6 minutes of data, for each pick-up. In order to support access by multiple clients simultaneously, new data is written into the buffer continuously and data older than 6 minutes is discarded. It is important to note that, as the achievable readout speed of the buffers is limited by the throughput of the Technical Network (1 Gbps), users can only extract a sub-set of the data from the last 6 minutes. Initial scripts for data readout have been developed by BE-ABP and the system has already been used extensively for MD purposes, including measurements of the ground vibration at IP1 [14]. An example of bunch-by-bunch injection oscillations for a complete batch shown in Fig. 7.

For 2016 it is planned to add dedicated "event" buffers that will automatically copy a snapshot of data for a certain number of turns around particular machine events (injection, beam dump, etc). The data can then be read from these buffers and processed offline. It is initially foreseen to include buffers for injection oscillations, instability triggers and post-mortem. Additional buffers can be added as required for operational or MD users.

As well as providing raw data, it is planned that the Observation Box will be used to perform online analysis. An initial application is the analysis of injection oscillations to dete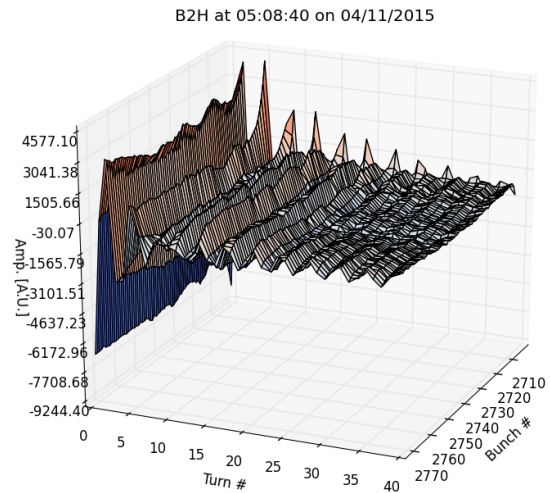rmine bunch-by-bunch parameters such as the ADT damping time and fractional tune [15]. The option of online, bunch-by-bunch, tune measurement throughout the fill has also been studied [16] and could be implemented at a later stage.

DATA STORAGE AND ANALYSIS

In the event of an instability trigger, a large amount of data needs to be analysed and potentially stored to allow experts to determine the cause of the instability. For example, the current LHC Head-Tail monitors can store 40 MB per trigger and the ADT Observation box can store 4 GB per trigger (assuming a 64 k turn buffer for all pick-ups). New digitizers with larger memories are currently being tested in the SPS and could also result in multi-GB acquisitions from the LHC Head-Tail monitors.

During 2015, BE-CO provided storage space on their operational NFS file servers for Head-Tail data. During the 2015 run approximately 1 TB of data was stored. Many of these files do not contain interesting data, but they were stored as identifying and removing them was a manual process. For 2016, the implementation of an automatic filter to remove uninteresting files will be studied. It has also been proposed to purchase a dedicated NFS server to avoid a situation where Head-Tail data fills the disks of the operational server and blocks other systems.

With the full deployment of the ADT Observation Box, online analysis of the data will be essential to only store the data corresponding to a real instability. Ultimately, something similar to the existing post-mortem system will be required to gather the information from many instruments in a single interface for online analysis by operators and specialists. Some very preliminary discussions have taken place during 2015, and will be continued in 2016.

CONCLUSION

The instrumentation and infrastructure related to instability diagnostics has seen major upgrades since run 1. The deployment of the LHC Instability Trigger Network allows triggers to be passed between instruments capable of detecting and diagnosing transverse instabilities. An initial detection algorithm has been implemented on the BBQ and this has been used to trigger the Head-Tail monitors throughout the year with good results. The recent deployment of the Observation Box will allow very large amounts of bunch-by-bunch data to be saved from the ADT. In order to fully exploit the huge amount of interesting data, it is clear that new frameworks for data storage and analysis will be needed for the upcoming runs.

REFERENCES

[1] E. Métral *et al*, "Review of the Instabilities Observed During the 2012 Run and Actions Taken" LHC Beam Operation Workshop, Evian, 2012.

[2] R.J. Steinhagen & W. Höfle, "Review on Functional Requirements on LHC Transverse Instability Diagnostics after LS1", CERN, Geneva, 2013.
http://indico.cern.ch/event/234589/

[3] J. Serrano *et al*, "The White Rabbit Project", IBIC'13, Oxford, 2013.

[4] T. Włostowski *et al*, "Trigger and RF Distribution Using White Rabbit", ICALEPCS'15, Melbourne, Australia, 2015.

[5] M. Cattin *et al*, "CERN's FMC Kit", ICALEPCS'13, San Francisco, 2013.

[6] S. Fartoukh & 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", LHC Project Report 602, CERN, Geneva, 2003.

[7] J. Ellis, "Level-1 Trigger Development for the Multi-Band Instability Monitor", CERN-STUDENTS-Note-2013-103, CERN, Geneva, 2013.

[8] R.J. Steinhagen *et al*, "A Multiband-Instability-Monitor for High-Frequency Intra-Bunch Beam Diagnostics" IBIC'13, Oxford, 2013.

[9] G. Kotzian *et al*, "LHC Transverse Feedback", *these proceedings*.

[10] M. Ojeda Sandonís *et al*, "Processing High-Bandwidth Bunch-by-Bunch Observation Data from the RF and Transverse Damper Systems of the LHC", ICALEPCS'15, Melbourne, 2015.

[11] Open Hardware Repository, "Simple PCIe FMC Carrier",
http://www.ohwr.org/projects/spec

[12] H. Timko *et al*, "Operational and Beam Dynamics Aspects of the RF System in 2015", *these proceedings*.

[13] W. Höfle, "Update on issues and plans with transverse instability monitoring", LHC Machine Committee 239, CERN, Geneva, 2015.
http://indico.cern.ch/event/453412/

[14] L.R. Carver, "First observations with the ADT observation box", LHC Beam Operation Committee 52, CERN, Geneva, 2015.
https://indico.cern.ch/event/461647/

[15] G. Kotzian & D. Valuch, "ADT parameter extraction: an alternative approach", LHC Beam Operation Committee 53, CERN, Geneva, 2015.
http://indico.cern.ch/event/467428/

[16] F. Dubouchet *et al*, "Tune Measurement from Transverse Feedback Signals in LHC", IBIC'13, Oxford, 2013.