WALL-CURRENT-MONITOR BASED GHOST AND SATELLITE BUNCH DETECTION IN THE CERN PS AND LHC ACCELERATORS

R. J. Steinhagen, S. Bart-Pedersen, J. Belleman, T. Bohl, H. Damerau, CERN, Geneva, Switzerland

Abstract

While most LHC detectors and instrumentation systems are optimised for a nominal bunch spacing of 25 ns, the LHC RF cavities themselves operate at the 10th harmonic of the maximum bunch frequency. Due to the beam production scheme and transfers in the injector chain, part of the nominally 'empty' RF buckets may contain particles, referred to as ghost or satellite bunches. These populations must be accurately quantified for high-precision experiments, luminosity calibration and control of parasitic particle encounters at the four LHC interaction points. This contribution summarises the wall-current-monitor based ghost and satellite bunch measurements in CERN's PS and LHC accelerators. Instrumentation set-up, post-processing and achieved performance are discussed.

INTRODUCTION

The Large Hadron Collide (LHC) is designed for a nominal bunch repetition frequency of 40 MHz while the RF cavities themself typically operate at a higher harmonic, implying that some of the buckets remain empty [1, 2, 3]. Due to imperfections, some of these buckets may be filled with minute amounts of particles, referred to as 'satellites' for percent-level filled buckets in the vicinity of nominal bunches, and 'ghosts' bunches in buckets that re-captured coasting particles during subsequent injection and with relative intensities of the order of 10^{-4}...10^{-6}-level. These unintentionally filled buckets need to be quantified as they generate additional collisions inside the LHC experiments before and after the nominal interaction points that may perturb precision measurements.

Being also detected by the LHC experiments and other devices such as the photon-counting-based longitudinal density monitor [4], these measurements involve long integration periods and colliding beams. The aim of this study is to assess the possibility whether unwanted bunch populations could be detected at an earlier stage, on much faster time-scales and exploiting the already existing wall-current monitor (WCM) infrastructure to allow mitigations already when they are being produced in the injector chain, notably in the Proton Synchrotron (PS) and Super-Proton Synchrotron (SPS).

EXPERIMENTAL SETUP

The WCMs used in the PS, SPS and LHC are based on an early 1970s design [5, 6, 7], consisting of a coaxial line with a small gap bridged by a defined impedance and surrounded by a closed vacuum chamber. This chamber acts as a DC shunt path, shields the gap from external fields and is loaded with ferrites to lower the cut-off frequency and to absorb the power entering the sump. The signal is extracted via eight SMA feed-through and re-combined using a star-topology combiner [8]. An effective combined frequency response from approximately 100 kHz up to 3-4 GHz is typically achieved. Minimising transmission losses and subsequent signal deformations, the combined signals are routed using short ($\approx 30\,\mathrm{m}$) 7/8" corrugated, high-bandwidth, coaxial PE-foam cable to a fast multi-GHz bandwidth oscilloscope, that is located in the closest possible, non-radiation environment in the accelerator tunnel.

To first order, the WCM output signal is an AC-coupled version of the longitudinal charge density, with the initial pulse shape being very short and determined by the detector's high-frequency response, and subsequent long-lasting undershoot being determined by the low cut-off frequency. As the droop amplitude is only a small $\approx 10^{-3}$ fraction of the initial peak amplitude for short LHC-type bunches, the integral over the initial pulse is a fair estimate of the bunch intensity. An exemplary signal response of the initial pulse is shown in Figure 1.

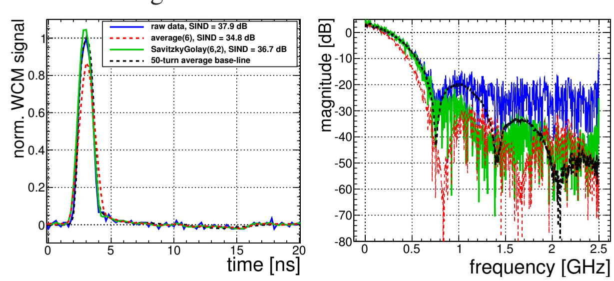

Figure 1: Time- and Fourier-domain comparison of the raw (blue), moving-average (red) and Savitzky-Golay (green) smoothing algorithm. The 50 turns averaged WCM signal (black) is indicated as reference.

MEASUREMENT PRINCIPLE

At first glance, the detection of the percent-level satellites appears to be limited rather by the digitiser's ADC quantisation than the WCM itself, which from a resolution point of view is limited only by imperfections and ultimately the thermal noise floor of the impedances and termination charges connected to it.

A range of acquisition systems with analog bandwidths above 2 GHz have been tested with similar analog performance figures: The measured signal-to-noise-and-distortion (SINAD) ratios[1] are typically in the order of about 44 dB, implying an approximate 1% accuracy over the ADC range. Further, the power spectrum of the residual

[1] measured as the ratio of the fundamental harmonic of a pure sinusoidal calibration standard to the first dominant spurious harmonic

ISBN 978-3-95450-121-2

noise floor (excluding the harmonics due to non-linearities) was found to be sufficiently flat (aka. 'white') within the targeted frequency range, which is a pre-requisite to allow improving the S/N-ratio by averaging successive turns [2].

Assuming a fixed sampling frequency of 10 (20) GS/s and typically short r.m.s. bunch lengths of about 0.25 ns in the SPS (LHC), the bunches are typically sampled by at least $n_s \approx 15(30)$ samples per bunch. Theoretically, with the given number of effective-number-of-bits (ENOB), this by itself yields already a relative bunch intensity resolution of about $1/\sqrt{n_s} \cdot 2^{ENOB} \approx 2 \cdot 10^{-3}$, but can be further improved by $\sqrt{n_{turn}}$ using turn-by-turn averaging. This averaging is limited essentially only by the required measurement bandwidth and time-scale of charge distribution and (to a lesser extend) by the WCM transfer function change. For practical purposes the limit is about 50 to 100 turns for the PS, given by the length of the stable period before the beams are extracted. Due to the significantly relaxed time-scales, the averaging in the LHC could be done up to about $112 \cdot 10^3$ turns, corresponding to about 10 seconds, but is presently limited to about 500 turns per ten seconds due to intrinsic acquisition hardware limitations (data transfer speed for the required 100 us sampling buffer). An upgrade is being investigated, which will improve the 0.5% duty cycle and which also allows some of the WCM response function compensation to be done by an FPGA.

Nevertheless, the achieved averaging performance is already quite adequate and yields resolutions of about 10^{-4} at 0.1 Hz measurement bandwidths, as for example shown in Figure 3.

SIGNAL POST-PROCESSING

By itself, the achieved resolution would already allow the detection of satellite and ghost bunch population 'by-eye', without further compensation. Albeit, further post-processing is required to compensate for signal deformations of the acquired raw signal to exploit this resolution also operationally for quantitative cycle-to-cycle (LHC: fill-to-fill) accelerator optimisations.

More precise estimates measuring bunch populations below the 10^{-3}-level require the compensation of the non-linear phase-delays, signal attenuation and recovery of the zero baseline, particularly if several bunches are circulating. As illustrated in Figure 2, one possible, fairly robust and flexible algorithm sequence was found to be: classical Wiener-deconvolution of the system response, followed by high-frequency noise rejection using polynomial regression [9], and base-line restoration using the SNIP background estimate [10, 11]. Subsequently, the satellite or ghost bunches can be detected with standard peak-detection routines discussed elsewhere.

The Wiener-deconvolution filter is based on the measured lab or beam-derived system response and sometimes

[2] Averaging over non-white-noise distributions such as '1/f' or Brownian-noise creates 'random-walk' effects that in time create artifical signal patterns that can be falsely interpreted as ghost or satellite bunch

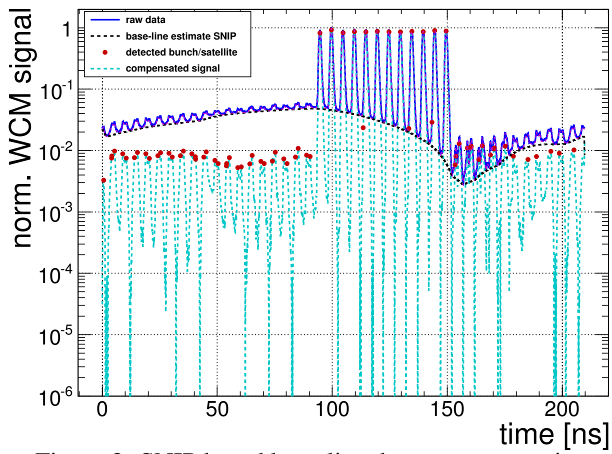

Figure 2: SNIP based base-line droop compensation.

complemented by an additional zero-pole filter that numerically shifts the lower cut-off frequency of the droop closer to zero. Its the preferred option for the measurement in the PS, as there is no high-precision numerical data available prior to its installation. The response function is obtained through the measured combined beam-based pick-up and cabling response to a well-defined single short bunch signal.

Low-pass filters were initially applied as part of the Wiener-deconvolution to improve the sample-to-sample noise particularly for higher-frequencies where the S/N-ratio is poor, but due to dynamic bunch shape and length changes, reduces the noise at the expense of the signal fidelity, i.e. perturbing the statistical features of the bunch distribution such as peak amplitude, bunch shape and length. The polynomial linear-regression filter described in [9] gives better results, as visible in the time- and Fourier-domain data shown in Figure 1.

Aiming at precise measurement of ghost and satellite bunch populations, it was noticed that the response function measurements vary on the sub-percent-level with time and apparently depend to second order also on the bunch filling pattern and signal amplitudes. Part of this effect is believed to be intrinsic to the properties of the dielectric and ferrite materials used in the pick-up, feed-throughs, combiner and cables, that may depend on atmospheric pressure, temperature, peak signal voltage and radiation effects.

These non-static imperfections affecting the droop of the signal, often much larger and hiding the to-be-detected ghost and satellite bunch signals, can effectively be compensated using the non-linear SNIP algorithm [10, 11] that requires only the size of the largest expected non-background structure (e.g. RF bucket width) as free parameter. Compensation examples are shown in Figure 3.

The Wiener-deconvolution has been omitted for better illustration of the background removal performance, and also because the lab- or beam-based system transfer function was not (yet) available for the PS installation. This omission causes some of the uncompensated reflections after the nominal bunches being erroneously identified as satellites. The achieved noise floor is about 10^{-4} for the

ISBN 978-3-95450-121-2

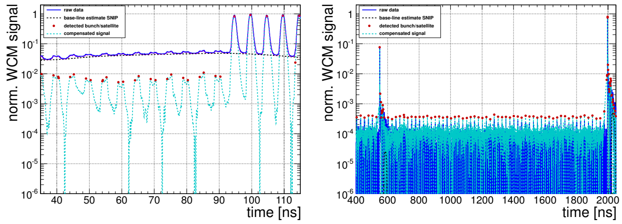

(a) PS 50 ns satellite example (b) LHC 25 ns satellite example

Figure 3: CPS and LHC satellite detection example showing the raw (blue), base-line estimate (black), compensated WCM (turquoise) and detected bunch and satellite peaks (red marker). The LHC example has been zoomed around the pilot and first nominal bunch. The 25 ns-spaced satellites are clearly visible.

PS and about 10^{-5} for the LHC detection scheme as expected from the statistical noise reduction prediction. The base-line is visibly flat which aides the subsequent bunch detection and intensity measurements.

DUAL-RANGE ACQUISITION

An alternative approach to achieve sub-percent resolutions within a turn via splitting and processing the signal in two amplitude ranges. One signal is split to one channel measuring the full range (and possibly applying the signal treatment as described above) and to a channel that is deliberately amplified and saturated for the nominal bunch signal to increase the sensitivity for satellite and ghost bunches. The tested oscilloscopes 'gracefully' clamp the signal with saturation recovery times below the nanosecond level. For systems that are not over-voltage protected, the clamping circuit illustrated in Figure 4 may be used. The additional attenuators after the splitters serve two purposes: to reduce the typical ample WCM signal from a few hundred volts to below about ± 5 V, and to improve the insulation between the ports, particularly to suppress the reflections in the regular channel created by the saturation and/or voltage clamping circuit.

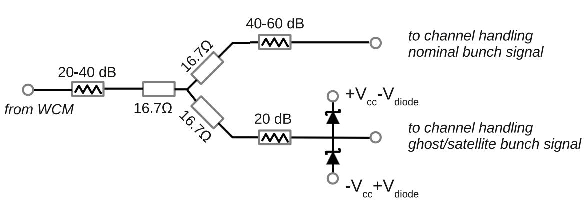

Figure 4: Dual-range signal clamping scheme.

Figure 5 shows the raw and compensated signals for a cycle in the PS believed to be satellite-free and one cycle with satellites being created deliberately at a known level. Despite the clamping, thwarting pure linear e.g. Wiener-deconvolution or zero-pole compensation of the droop, the discussed high-frequency noise rejection and background removal algorithms work very well. The achieved resolution performance of this method despite using only one turn is about 10^{-5}, which is very appealing and allows the satellite detection in cases that were previously inaccessible.

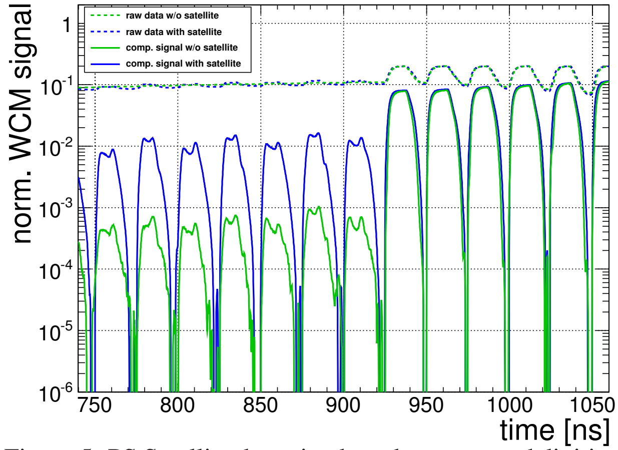

Figure 5: PS Satellite detection based on saturated digitiser inputs. The raw data (dashed lines) and reference with reduced (green) and enhanced satellites (blue) are indicated. The vertical scale has been re-scaled to correspond to the full scale indicated in Figure 2.

CONCLUSIONS

The LHC and its injectors operate with filling patterns containing many empty RF buckets, which due to imperfections in the beam production may be filled with minute amounts of particles referred to as satellite and ghost bunch populations. These populations can be precisely measured using the existing WCM installations and turn-by-turn averaging of the repetitive bunch filling pattern, or saturation of the digitiser's input channel. A relative bucket-by-bucket intensity resolutions down to the 10^{-5}-level can be achieved. In order to exploit this resolution, the pick-up and cabling imperfections require a precise compensation for an automatic detection and quantitative assessment of these ghost and satellite populations. This can be achieved through a combined Wiener-deconvolution, polynomial linear-regression and SNIP-based background removal. While the percent-level compensation depends on the measure lab and beam-based calibration data, the sub-percent response is compensated independently on the knowledge on the pick-up response and automatically adapted to the given input beam signal.

REFERENCES

[1] "LHC Design Report, CERN-2004-003, CERN, 2004.
[2] M. Benedikt (ed.), CERN-2000-003, CERN, 2000.
[3] R. Steerenberg et al., Chamonix 2011, CERN, 2011.
[4] A. Jeff et al., NIMA, Vol. 659, Issue 1, 2011, pp. 549-556.
[5] T. Linnecar, CERN-SPS/ARF/78-17, 1978.
[6] T. Bohl, CERN-BE-2009-006, 2009.
[7] R. Cappi et al., CERN-PS-87-31-PSR, PAC 1987, 1987.
[8] J. Durand, PS/LP/Note 94-14 (Tech.), CERN, 1994.
[9] A. Savitzky & M. Golay, Anal. Chem., Vol. 36, No. 8, 1964.
[10] C.G. Ryan et al., NIM B34 (1988), 396-402.
[11] M. Morháč et al., NIM, A401 (1997) 113-132.

ISBN 978-3-95450-121-2