TUNE MEASUREMENT FROM TRANSVERSE FEEDBACK SIGNALS IN LHC

F. Dubouchet*, W. Höfle, G. Kotzian, T.E. Levens, D. Valuch, CERN, Geneva, Switzerland
P. Albuquerque, HEPIA, Geneva, Switzerland

Abstract

We show how bunch-by-bunch position data from the LHC transverse feedback system can be used to determine the transverse tunes. Results from machine development experiments are presented and compared with theoretical predictions. In the absence of external beam excitations the tune is visible in the spectra of the position data with the feedback loop as a dip, while with external excitation a peak is visible. Both options, observation with and without excitation, are demonstrated to be complementary. Periodic excitation and observation of the free oscillation can also be used to determine the damping time of the feedback in addition to the coherent tune. Plans are outlined for hardware upgrades of the LHC transverse feedback system that will enable fast online processing of bunch-by-bunch, turn-by-turn data using Graphical Processing Units (GPUs). By using GPUs we gain the ability to compute and store the spectrum of all bunches in real-time and the possibility to reconfigure, test and deploy algorithms. This data acquisition and analysis architecture also allows changes to be made without disturbing the operation.

INTRODUCTION

Maintaining the correct tune throughout acceleration is of paramount importance for the Large Hadron Collider (LHC) at CERN. The procedure to adjust the tune and chromaticity prior to the injection of a high intensity beam for physics involves injection of a low intensity pilot bunch of between 5×10^9 and 1×10^{10} protons. Residual oscillations of this bunch are detected by the BBQ system [1] which is based on diode detection. For the low intensity bunches that are used to adjust the machine, the transverse feedback system is not needed and the measurement of the tune using the BBQ system is both reliable and precise [2]. For the operational beam, used for physics, which has had up to $\simeq 1.6\times10^{11}$ protons per bunch and 1380 bunches per beam during the first LHC proton run (2010-2012), transverse feedback is needed throughout the cycle to stabilize the beam against coupled bunch instabilities [3, 4]. The gain of this feedback is high with damping times of the order of 20 turns, suppressing not only instabilities but also reducing emittance blow-up by rapidly damping injection errors and reducing the detrimental effect due to external perturbations, such as noise on magnets that shake the beam. Consequently, the residual oscillations of the beam are reduced as desired but, together with additional broadband noise injected by the feedback system itself, the signal-to-noise ration of the BBQ based tune measurement is greatly reduced. During the ramp a tune feedback is needed and the quality of the BBQ based signal has been a concern, as an incorrect measurement of the tune can lead to wrong adjustment of the tune quadrupoles effectively driving the tune away from the desired value and leading to beam loss.

A number of options to improve the tune measurement have been explored. A method that could be implemented immediately was the reduction of the transverse feedback gain for a single batch in the beam and a gating of the BBQ signal to this particular bunch train. Standard operation of the LHC in 2012 used this method which has improved the reliability of the tune measurement and tune feedback. In some cases, the oscillations have been enhanced by sweeping an excitation signal across the tune with the transverse feedback kickers.

In the following we report on alternative methods of tune measurement that use the digitized bunch-by-bunch oscillation data present in the transverse feedback loop itself to derive a reliable tune measurement. With a set of eight pick-ups, two per plane (horizontal and vertical) per beam, generating data of 16 bit each at the bunch repetition rate of 40 MHz, the total data stream present in the transverse feedback system amounts to over 5 Gbit/s. The transverse feedback system is a VME based and, due to limited bus bandwidth, only allows a fraction of the data to be extracted for analysis. The data stored during machine development session (MDs) is being used for offline analysis to guide the development more universal system based on Graphics Processing Units (GPUs) that will receive the real-time data via a serial link for online analysis in order to display the tune and for the tune feedback system. It has been important to test the feasibility of different algorithms with real data and on different GPUs in order to select the correct technologies for the planned system.

In general two options for tune measurement using the transverse feedback data were tested: firstly a purely passive observation of the data stream and secondly active excitation on selected bunches using the transverse feedback system to create measurable coherent oscillations. The latter can also be used to monitor the damping, however this will lead to some blow-up of the selected bunches and is thus invasive.

MEASUREMENT PRINCIPLES

In a continuous time model, the transverse betatron oscillations of the beam and the active damping by the feed-

* frederic.dubouchet@cern.ch

ISBN 978-3-95450-127-4

back system can be described by a second order differential equation whereby the machine tune and damping terms are determined by the complex eigen frequencies. A Fourier analysis of the time domain data sampled at every turn yields the tune of the beam and different regimes can be distinguished. With active feedback, oscillations are driven by residual noise on magnets and a tune line is visible, but this is not very reliable as its strength is governed by an uncontrolled excitation. Active excitation can be used to drive the beam and make the tune line grow above the measurement background noise, but usually this is at the expense of emittance blow-up. In the following we describe the gated periodic excitation that has been developed for, and tested in, the LHC during 2012. Passive observations can also be made without any additional excitation. With feedback on a “trench” at the tune develops and we have to search for a minimum in order to determine this value. This technique has been used in electron storage rings [5] and a detailed analysis using z-transforms has been presented in [6]. Note that this technique can distinguish very well between the external lines in the spectrum at multiples of 50 or 100 Hz, observed particularly in the horizontal plane in the LHC, and the real tune value.

RESULTS AND DISCUSSION

Tune Measurement with Selective Excitation

In the case of a tune measurement with selective excitation, the transverse feedback kickers are used to excite the leading batch of the beam. For the standard 50 ns LHC beam, the first high intensity batch in every physics fill consists of 6 bunches and the applied excitation was gated to these. As a single turn kick by the transverse feedback results only in a small oscillation amplitude, a series of kicks is applied on successive turns, the amplitudes of which is modulated at the nominal tune frequency. An excitation lasting for around four to five turns is sufficient to generate an oscillation of approximately 100 μm at a $\beta \simeq 150 - 200\ m$ and this can be detected with good precision by the transverse feedback system. The *free* oscillation following the excitation is recorded and the analysis only considers the data from this free oscillation (see figure 1). The DFT of the data stream readily gives the tune, figure 2 and, in the presence of active damping by the feedback, also the damping time.

In practice this method suffered from difficulties associated with the presence of spurious signals in the spectrum. The data taken for 6 bunches (2730 turns per bunch) in the presence of strong damping produces a wide tune signal with some internal structure but, in the presence of such strong damping, the different lines could not be resolved. However, looking for the peak of the signal can generate an output tune signal that jumps between different values (lines) of tune, this can be seen on figure 3. It is not clear if the different lines that are visible are artefacts, externally driven, or a true representation of the beam’s eigen frequencies. All six bunches exhibited the same behaviour.

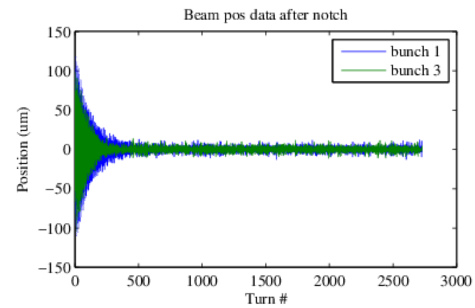

Figure 1: Example of Tune Measurement with Selective Excitation (time domain signal).

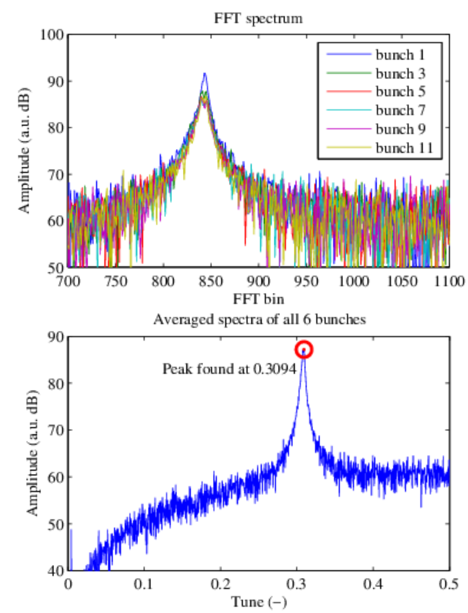

Figure 2: Spectra for individual bunches (top) and averaged spectra over 6 bunches (bottom).

Two possible remedies have been proposed. Firstly, the reduction of the feedback gain for the leading six bunches, which was prepared in 2012, but has not yet tested with beam for this purpose. Secondly an analysis of the oscillations which will integrate out the different lines present in the tune spectrum. The first approach will give more insight in the structure of the tune spectrum and the origin of the spurious lines, while the second approach can be used for a robust tune measurement in view of using the result for the tune feedback. A means of integrating over the oscillations is to look at the average phase advance over a given number of turns using Hilbert transforms of the time domain data.

Non-Invasive Tune Measurement from Transverse Feedback

As mentioned in the previous chapter, when there is a functioning transverse feedback, the tune observation can

ISBN 978-3-95450-127-4

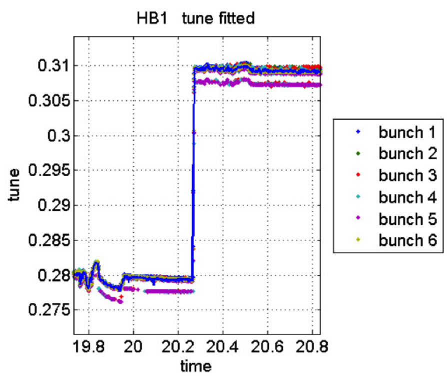

Figure 3: Example of Tune Measurement (peak DFT) for six bunches of beam 1 in the horizontal plane during the ramp and squeeze with two lines spaced at 20 Hz. During the squeeze the fractional tune changes from 0.28 to 0.31.

be made without additional active excitation of the beam. The tune value can be observed as a trench (or a "dip") in the in-loop signal spectrum. Provided that the in-loop data is not too noisy, and that the record length is sufficient, the tune can be computed as a simple FFT for every individual bunch. However, in the real system, the noise floor of a typical single acquisition (1 bunch, 2048 turns) is too high to obtain any meaningful reading of the tune. The transverse damper in LHC treats every bunch individually, and the in-loop data is available at the bunch-crossing frequency of 40 MHz. Providing that there is sufficient computing power, it is possible to acquire 2048 turns of data for all of the circulating bunches (e.g. 1380 in 50 ns operation, or 2808 in 25 ns operation), compute the individual FFTs and average the magnitudes of the spectra. This will lower the noise background and allow us to observe a clear trench at the tune value.

This technique has been demonstrated for the first train of 6 bunches during normal LHC operation. The data has been collected for several minutes to accumulate a sufficient number of datasets for averaging and the data has later been processed offline. Figure 4 and 5 shows the results of the non-invasive tune measurement with an accumulation value of 10 (i.e. 60 spectra averaged). The tune trench is clearly visible. The resolution will significantly increase if we average the spectra of all bunches, whilst still obtaining a throughput of 5 tune values per second.

TEST OF ALGORITHMS ON A GPU USING EXPERIMENTAL DATA

With data acquired during MDs, we have been able to build an off-line analysis tool [7] and try different algorithms. To compute the tune in real time the data must pass through the different steps within certain time constraints. The whole system has to be able to process all of the bunches of the machine, up to a maximum of 2808, at a rate faster than the acquisition rate. The idea is to produce 2048 point Fast Fourier Transforms (FFTs) faster than the

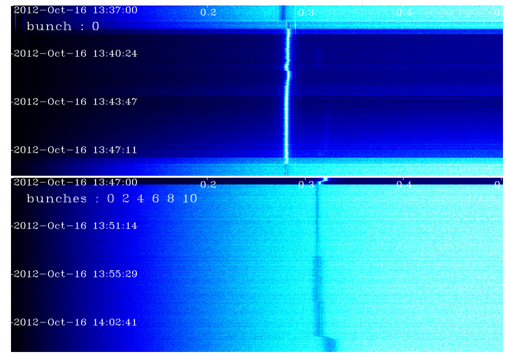

Figure 4: Tune Measurement (non-invasive), with low damper gain (top) and with high damper gain (bottom).

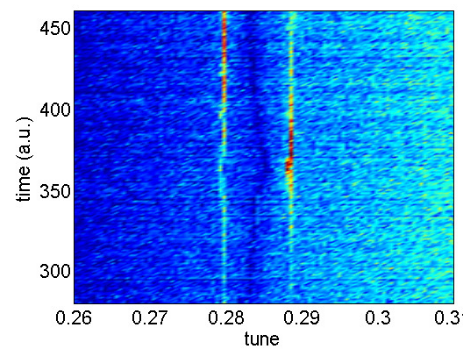

Figure 5: Detail of Tune Measurement (non-invasive) with spurious lines spaced by 100 Hz and tune (trench) in between.

speed of acquisition. Using conventional central processing units (CPUs), as shown on Table 1, we are slower than the acquisition speed of around 100 ms.

Table 1: Speed for 3000 Acquisitions of 2048 Points Comparing Threads, Frequencies and Different Models of CPUs and a GPU

Device	Type	Threads	[GHz]	[ms]
i7-3720QM	FFTW	8	2.6	310
i7-3720QM	OpenCL	8	2.6	272
Tesla M2090	OpenCL	512	1.3	35

A modular software design has allowed testing of an implementation using a mixture of processing between CPUs and GPUs. OpenCL [8], the open standard for general purpose GPU programming, allows us to have the same code running on both the CPU and the GPU. To have an external comparison, the library FFTW [9] has been used and this exhibits a very similar level of performance as OpenCL on the CPU. On the GPU, computing is about ten times faster than on the CPU, even when considering a GPU model that is a generation old. This has showed that we can compute the tune FFTs for each individual bunch faster than the acquisition speed. The time distribution of the different parts of the algorithm is displayed in figure 6.

ISBN 978-3-95450-127-4

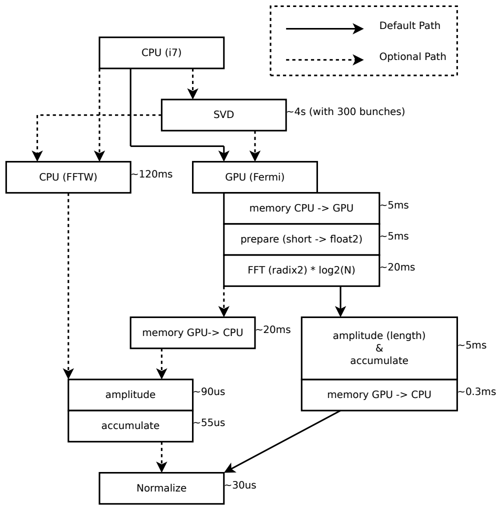

Figure 6: Time flow with different implementations and with 3000 bunches of 2048 points each.

Parallelization of FFTs on GPUs can be done by distributing all of the twiddling and multiplication steps between the available cores so that, in our case, we have a maximum possible parallelization optimization of 2880 times 1024 which is three order of magnitude larger than the number of core on the tested GPU. Therefore, by utilising GPUs with more cores, we should be able substantially reduce the processing time through increased parallelisation. Acquisition and computing speed show that it is possible to build a real-time bunch-by-bunch tune acquisition system. This system should allow us better diagnostics in the LHC and allow a smooth transition to higher energy after the current long shutdown of the LHC (LS1).

ARCHITECTURE FOR A FUTURE ONLINE OPERATIONAL SYSTEM

In the transverse feedback system we have the digitized radial position data available from each pick-up as a serial stream at a bunch by bunch rate of 40 MHz. This results a 1 Gbps serial line rate after 8b10b encoding and comma padding. We plan to perform pre-processing of the data streams in the ADT hardware in order to combine all pick-ups and produce a single 1 Gbps stream of data for each plane. In order to decouple the tune measurement and observation from the operation of the transverse feedback system, we will stream this data out of the feedback system over fiber to a dedicated sampling box based on a PC platform. A dedicated fiber to PCI Express interface, based on the open-hardware "SPEC" card developed by BE-CO [10], will capture the serial stream into an on-board memory buffer and make it available to the CPU over the PCI Express bus. The serial stream can easily be replicated using fiber splitters and sent to a virtually unlimited number of sampling boxes which could be dedicated to a particular user (e.g. operations group), or a particular analysis (tune measurement, instability trigger, various online bunch by bunch analysis). By selecting an external solution, based on off the shelf PC hardware, we can easily upgrade the sampling boxes provide extra capacity in terms of storage space and GPU processing power without disrupting the operation of the transverse feedback system.

CONCLUSIONS

In machine developments experiments in the LHC, we have tested two methods to derive the transverse tunes from the feedback data. The first method uses a several turn excitation burst that is gated to a short train of bunches and the second, non-invasive, method consists of observing the feedback signal of the individual bunches without external excitation. Both methods produced promising results. Algorithms to calculate the DFTs required on GPUs have been tested. The development of a hardware platform based on GPUs is proposed to receive the streamed data from the transverse feedback system for on-line processing. This hardware platform can process the data stream without obstructing communications with the damper feedback hardware which is realized in VME standard.

ACKNOWLEDGMENT

We would like to acknowledge the help of the BE-OP group at CERN for the support given in carrying out the experiments in the LHC and for setting-up the beams.

REFERENCES

[1] M. Gasior, DIPAC'11, Hamburg, May 2011, pp. 445–447.

[2] R. J. Steinhagen, PAC'11, New York, Mar 2011, pp. 1399–1403.

[3] W. Höfle et al., EPAC'08, Genoa, May 2008, pp. 3266–3268.

[4] W. Höfle, IPAC'13, Shanghai, May 2013, pp. 3990–3994.

[5] C. H. Kuo et al., IPAC'11, San Sebastián, September 2011, pp. 514–516.

[6] C.-Y. Yao et al., IPAC'13, Shanghai, May 2013, pp. 3058–3060.

[7] F. Dubouchet, "Betatron tune measurement with the LHC damper using a GPU", CERN Thesis (unpublished), 2013.

[8] M. Scarpino, *OpenCL in Action*, (Manning Publications, 2011).

[9] M. Frigo, S.G. Johnson, "The Design and Implementation of FFTW3", Proceedings of the IEEE, February 2005, IEEE93, p. 216-231, 0018-9219

[10] `http://www.ohwr.org/projects/spec`

ISBN 978-3-95450-127-4