ARCHITECTURE OF TRANSVERSE MULTI-BUNCH FEEDBACK PROCESSOR AT DIAMOND

M.G. Abbott, G. Rehm, I.S. Uzun, Diamond Light Source, Oxfordshire, UK

Abstract

We describe the detailed internal architecture of the Transverse Multi-Bunch Feedback processor used at Diamond for control of multi-bunch instabilities and measurement of betatron tunes. Bunch by bunch selectable control over feedback filters, gain and excitation allows fine control over feedback, allowing for example the single bunch in a hybrid or camshaft fill pattern to be controlled independently from the bunch train. It is also possible to excite all bunches at a single frequency while simultaneously sweeping the excitation for tune measurement of a few selected bunches. The single frequency excitation has been used for continuous measurement of the beta-function. A simple programmable event sequencer provides support for up to 7 steps of programmable sweeps and changes to feedback and excitation, allowing a variety of complex and precisely timed beam characterisation experiments including grow-damp measurements in unstable conditions and programmed bunch cleaning. Finally input and output compensation filters allow for correction of front end and amplifier phasing at higher frequencies.

INTRODUCTION

At Diamond Light Source (DLS) we have been using a Transverse Multi-Bunch Feedback (TMBF) system based on the Libera hardware platform [1] for nearly a decade [2–5] to control multibunch instabilities. With a new bunch position every 2 ns it is appropriate for the feedback processing to be implemented in an FPGA. The original firmware was based on work done at the ESRF [6], but subsequently the firmware and control system have been rewritten to provide further extensions to the system functionality [7, 8]. This work has extended the lifetime of the original platform and increased its capability, and the DLS TMBF system has now been adopted at ALBA [9] and is being evaluated at the Australian Synchrotron.

In this paper we expand on and update the original IBIC 2013 [7] paper with many system refinements.

OVERVIEW

Figure 1 shows the TMBF processors in context. A 4-button electrode assembly together with an RF hybrid circuit picks up the horizontal and vertical position of each bunch. An RF front end converts this raw 3rd harmonic signal into a 250 MHz bandwidth bunch position signal suitable for processing by TMBF. This signal is digitised at 500 MHz (using four 125 MHz ADCs phased at 2 ns intervals), processed by the FPGA, and then reconverted to drive the striplines.

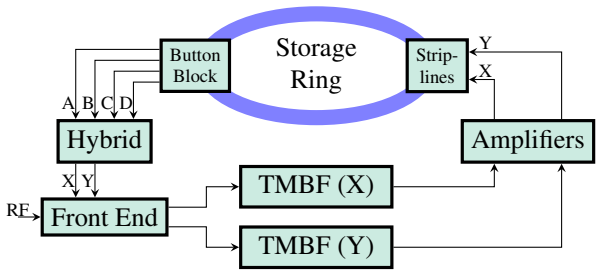

Figure 1: TMBF system in context. The Transverse Multi-Bunch Feedback system measures the position of each bunch, detects the betatron oscillations of each bunch, and generates a drive signal to suppress the oscillations.

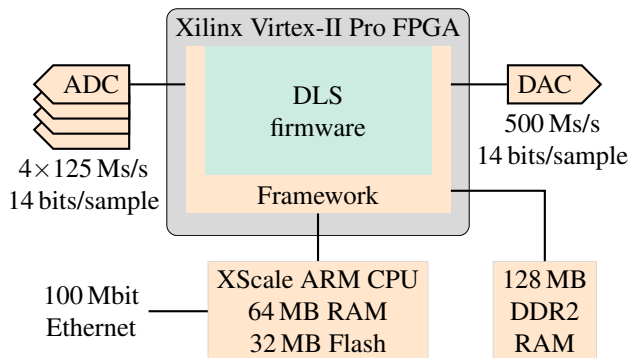

Figure 2: Libera System Platform. The DLS TMBF system is implemented on the Instrumentation Technologies Libera platform with the control system running EPICS on an ARM based embedded Single Board Controller (SBC).

The position of each bunch oscillates at the machine betatron tune, and these oscillations can normally be suppressed by feeding them back with a phase shift of 180°. The TMBF measures the horizontal or vertical position of each bunch, runs a 10-tap FIR filter on each bunch position to compute the appropriate phase adjustment, and outputs a negative feedback signal.

The TMBF processing system consists of analogue to/from digital converters connected to an FPGA for the high speed processing, together with memory for data capture and an embedded Single Board Computer as shown in Fig. 2.

As well as the core feedback function, a number of diagnostic functions are provided, both to enable easy monitoring of the system status, and to support complex experiments on the beam. System status monitoring includes detailed overflow detection and quick measurement of beam movement. Complex experiments are supported by internal oscillators, detectors, and a programmable sequencer, described in more detail below.

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

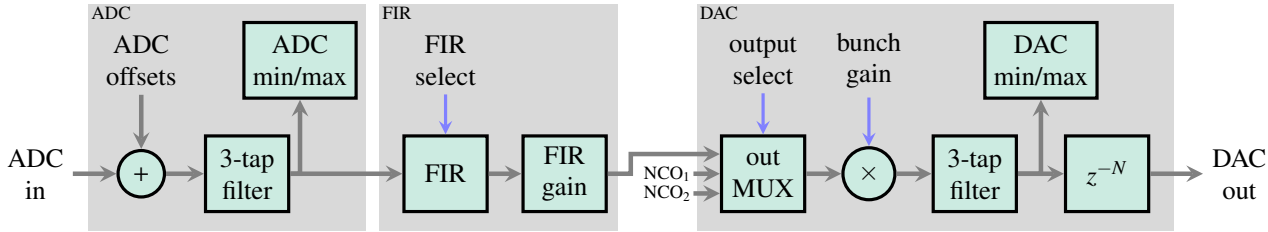

Figure 3: This is the core data processing chain in the TMBF processor. Data processing starts by adding a DC offset to each of the four ADC channels to compensate for static ADC errors, followed by a 3-tap filter to compensate for high frequency phase errors in the front end. The minimum and maximum value per bunch of both the ADC and DAC streams is captured for display. A 10-tap filter with programmable gain (in 6 dB steps) is applied in turn to each bunch in the ring. The output multiplexer adds any combination of its three inputs, which is then scaled by a bunch specific gain. Finally an output pre-emphasis filter corrects for amplifier errors and is followed by a delay line to correctly close the loop.

FEEDBACK DATA PROCESSING

The main function of TMBF is to stabilise transverse oscillations of the beam. This is done by running a separate 10-tap FIR on the position of each of the stored bunches. The core data processing chain shown in Fig. 3 combines this feedback with up to two optional Numerically Controlled Oscillator (NCO) outputs.

Note that numerical overflow can occur at any stage in this processing chain. When this is detected, a saturated output is generated and the resulting overflow event is reported to the user through the EPICS interface.

ADC Data In

The Libera platform uses four separate 125 MHz ADCs to sample the 500 MHz bunch position signal, and so the ADC input arrives as four parallel streams of 125 MHz data. The first processing step is to compensate for small DC offsets between the ADCs. These offsets are easily measured and compensated.

The next step is to compensate for gain differences and high frequency phase and gain errors in the front end. This is done by running a channel dependent 3-tap filter on the input data stream.

Finally a "min/max" subunit measures the minimum and maximum compensated value for each bunch; this is read out at 100 ms intervals and used to provide an accessible overview of bunch motion.

FIR Feedback Processing

Next a 10-tap FIR is run on the stream from each bunch. In practice, this means that 936 separate FIRs are run, with a delay of 936 bunches between each tap. The FIR unit can be programmed with four different filters, each of which can be separately selected for each bunch. The filter taps and final output gain are statically configured through the EPICS interface.

The ability to control bunches independently has two main uses: firstly, it can be used to apply a different filter to the isolated bunch in a hybrid fill; secondly, it can be useful as part of detailed machine physics experiments.

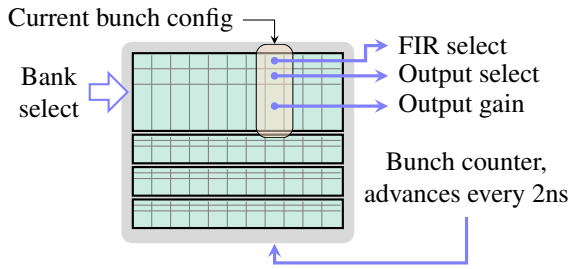

Figure 4: Bunch by bunch FIR and output control. The sequencer selects one of four banks to be active, defining basic machine behaviour. Each bank contains bunch configuration information for each of the stored bunches, used to control the feedback filter and the output configuration.

DAC Data Out

Finally the filtered signal is prepared for output. At this point three candidate output signals are selected and summed together: the FIR filtered signal and two internally generate NCO signals. Each of these three signals can be output or set to zero, and this control is per bunch.

Next the summed output is scaled by a bunch specific scaling factor. As this is signed it can be used to reverse the orientation of feedback on a single bunch.

Finally an output pre-emphasis filter is run to compensate for amplifier effects and the output is delayed so that the complete loop delay is a full machine revolution.

In the same way for ADC data, a "min/max" component measures the movement of each output bunch over a measurement interval of 100 ms, providing a quick view of bunch output motion.

Bunch by Bunch Control

As noted above, the FIR filter, the output sum, and the final output gain are selected per bunch. The "Bunch Select" unit, see Fig. 4, stores four arrays (or "banks") of bunch specific configurations — one configuration is selected by the sequencer unit (described below), and the current bunch configuration advances on each bunch. This configuration can be written through the EPICS interface.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors

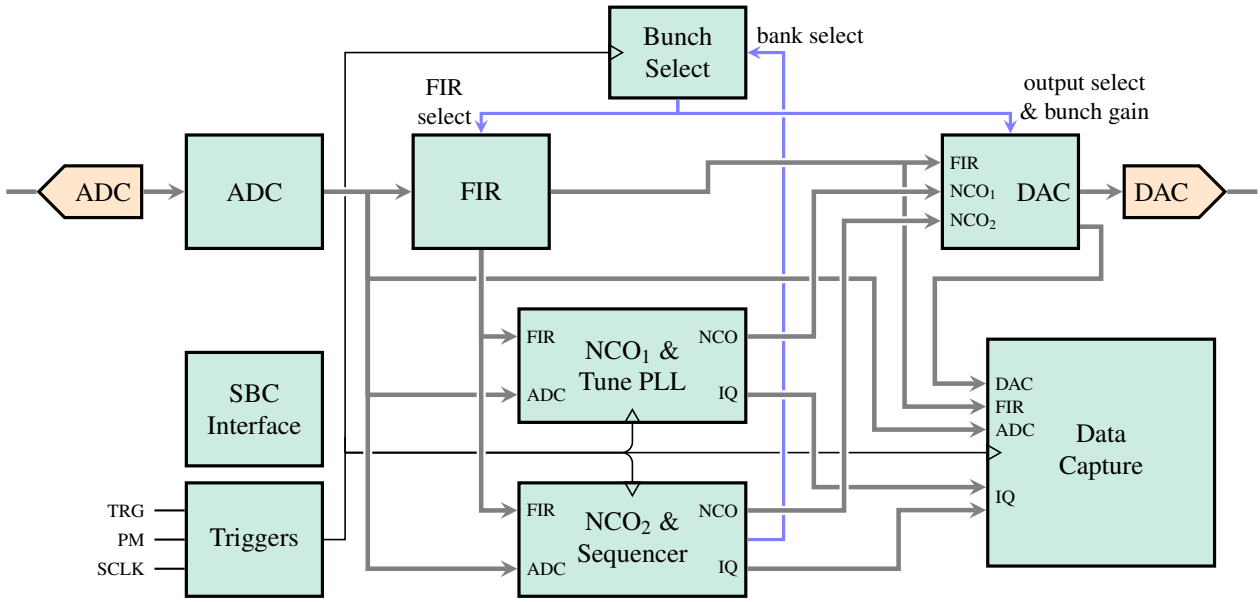

Figure 5: Here we see all of the major blocks of the FPGA system design and their data interconnections. FPGA blocks are shown thus: □ and analogue/digital converters thus: ▷. The main data flow is from the ADC, through the FIR with a separate FIR filter selected for each bunch, and out through the DAC with the option of adding up to two internally generated sine waves. The other paths are for control and data capture. The SBC interface controls and communicates with all other components of the system: the EPICS interface is through this component.

EXTRA DIAGNOSTIC DATA PROCESSING

Figure 5 shows the complete DLS TMBF firmware. The ADC-FIR-DAC chain has been described above. The remaining major components are for extra diagnostics and advanced machine experiments.

Data Capture

As already shown, min/max overview data is available for ADC input and DAC output. The Data Capture unit allows more detailed information: up to 4096 bunches of two out of three of DAC/FIR/ADC data streams can be captured, or up to 64 million bunches (more than 65,000 turns) of any one input can be captured. The shorter data capture is to FPGA block memory, the larger to external DDR2 RAM.

An alternative capture source is detector IQ data from either of the two NCO units. In particular, this is used for tune sweeps and more complex machine experiments.

NCO_1 and Tune PLL

Figure 6 shows the first NCO unit together with the tune tracking application (Tune PLL). When tune tracking is inactive the NCO generates a fixed frequency sine wave with programmable gain which can be used to drive selected bunches.

When tune tracking is active the detector (see Fig. 7 for the detector structure) measures the phase response of the beam at the currently driven frequency and runs a simple PI feedback loop to adjust the frequency to keep the phase response static. The frequency shift can be read as a continuous stream through the EPICS interface.

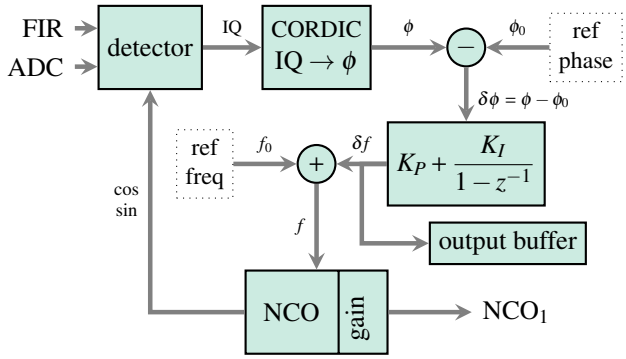

Figure 6: The oscillator NCO_1, can be set to a fixed frequency, or can be used as part of a tune tracking phase locked loop. The tune phase ϕ is measured at the operating frequency f and used to compute a sequence of frequency corrections δf to maintain the phase error $\delta\phi$ at zero. The reference phase and frequency are configured via EPICS.

This is very useful for measuring rapid changes in the machine betatron frequency, but will only work when the centre frequency and phase are already known to sufficient precision. Measuring this is the primary job of the next unit.

NCO_2 and the Sequencer

Figure 7 shows the second NCO unit together with the sequencer unit and a detailed view of the detector. The detector measures the complex response of the selected input to the driving frequency. The duration of each

Copyright © 2015 CC-BY-3.0 and by the respective authors

ISBN 978-3-95450-148-9

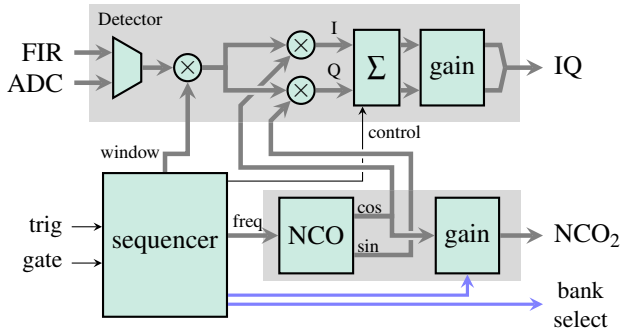

Figure 7: The oscillator NCO_2 is under control of the sequencer, and is only enabled while the sequencer is running a program. The sequencer controls the NCO frequency and gain and performs programmed frequency sweeps. Data from the beam is mixed with the excitation waveform in the detector to measure a complete complex IQ response. Operation of this system produces a waveform of IQ measurements which can be used for measuring betatron tune and other machine parameters.

measurement is controlled by the sequencer. The measured IQ value is scaled to a pair of 16-bit numbers for storage and transferred to the control system.

The TMBF sequencer controls the NCO_2 frequency, the NCO_2 output level, and which bunch bank is currently active. A machine experiment is performed by programming up to seven different control configurations and then triggering the start of the sequencer. The sequencer will then work its way through the programmed states, changing the output configuration as appropriate and capturing the appropriately programmed number of IQ samples, before finishing. Typically the sequencer and the data capture unit are triggered together.

For example, a standard tune sweep is performed using a single sequencer state which configures NCO_2 for output and steps through a preset range of NCO frequencies. Alternatively, a grow-damp experiment requires more states: one to excite the beam for a selected period, one to allow the beam to grow without feedback or excitation for a period, and one to restore feedback.

One final feature of the sequencer is the so-called "super-sequencer": this repeatedly performs a sequencer programme for a series of base frequencies. This is used to perform a grow-damp experiment on each transverse mode in turn, and allows the transverse mode damping times to be measured for all modes in a fraction of a second [8].

EPICS CONTROL SYSTEM

The control system running on the embedded ARM processor was developed concurrently with the FPGA firmware, and provides access to all settings and readings through around 560 EPICS PVs. A number of supporting scripts were written in Python and Matlab to help configure some of the more complex functions of TMBF.

A complete set of user interface screens built using EDM provides access to every PV published by TMBF. Around 255 PVs are used to configure settings, so the supporting scripts are important for normal operation.

The EPICS control system and the TMBF FPGA system are tightly coupled. Although the ARM processor is low powered and lacks hardware floating point support it is capable of some signal processing if care is taken. In particular, tune measurement is done by first performing a frequency response sweep with the FPGA, and then analysing the resulting IQ response and fitting a multi-pole tune response model. It has been possible to develop quite a sophisticated tune measurement process on this system.

CONCLUSIONS

The upgraded TMBF system described here has been in use at Diamond since late 2013 and is allowing us to perform very complex measurements to characterise the beam behaviour. This development has pushed the 10 year old FPGA technology to its limits, the FPGA now has little or no room for further developments.

Our next step at Diamond is to implement longitudinal bunch-by-bunch feedback in preparation for operation with normally conducting cavities. This will be based on a more modern FPGA on industry standard MicroTCA hardware.

REFERENCES

[1] Instrumentation Technologies, *Libera Bunch-by-Bunch*, `http://www.i-tech.si`.

[2] A.F.D. Morgan, G. Rehm, I. Uzun, *First Tests of the Transverse Multibunch Feedback at Diamond*, DIPAC 2007.

[3] A.F.D. Morgan, G. Rehm, I. Uzun, *Performance and Features of the Diamond TMBF System*, EPAC 2008.

[4] G. Rehm, M.G. Abbott, A.F.D. Morgan, J. Rowland, I. Uzun, *Measurement of Lattice Parameters Without Visible Disturbance to User Beam at Diamond Light Source*, BIW 2010.

[5] I. Uzun, M.G. Abbott, M.T. Heron, A.F.D. Morgan, G. Rehm, *Operational Status of the Transverse Multibunch Feedback System at Diamond*, ICALEPCS 2011.

[6] E. Plouviez, P. Arnoux, F. Epaud, J. Jacob, J.M. Koch, N. Michel, G.A. Naylor, J.-L. Revol, V. Serriere, D. Vial, *Broadband Bunch by Bunch Feedback for the ESRF using a Single High Resolution and Fast Sampling FPGA DSP*, EPAC 2006.

[7] M.G. Abbott, G. Rehm, I.S. Uzun, *Capability Upgrade of the Diamond Transverse Multibunch Feedback*, IBIC 2013.

[8] G. Rehm, M.G. Abbott, A.F.D Morgan, *New Features and Measurements using the Upgraded Transverse Multibunch Feedback at Diamond*, IBIC 2014.

[9] A. Olmos, U. Iriso, J. Moldes, F. Pérez, M. Abbott, G. Rehm, I. Uzun, *Integration of the Diamond Transverse Multibunch Feedback System at ALBA*, IBIC 2015.

ISBN 978-3-95450-148-9

Copyright © 2015 CC-BY-3.0 and by the respective authors