

FPGA BASED FAST ORBIT FEEDBACK SYSTEM FOR THE AUSTRALIAN SYNCHROTRON

Y-R. E. Tan, T. Cornall, S. A. Griffiths, S. Murphy, E. Vettoor,
Australian Synchrotron, Clayton, Australia

Abstract

An initial design for a Fast Global Orbit Feedback System based on FPGAs has been proposed for the Australian Synchrotron Light Source (ASLS). The design uses a central processor (Xilinx Virtex 6) for all the computations and fast optical connections to distribute the computed data to corrector magnet power supplies. The network topology consists of two fibre optic rings. The first ring is used by the Libera Electron's to aggregate the beam position data at 10 kHz using Instrumentation Technologies' Grouping algorithm. The second ring is used to transmit the computed data. The cycle frequency of the feedback is 10 kHz with a targeted total latency of under 400 μ s. We shall give an overview of the design goals and discuss the merits of the current implementation. We shall also present the measured bandwidth of the stainless steel vacuum chamber and test results from initial prototyping work.

INTRODUCTION

The ASLS is a 3 GeV light source open to users since 2007. Although not implemented at the start, the plan was to implement the fast global orbit feedback system at a later stage when the need arose. The beam stability requirement is that the RMS of the beam motion be < 10% of one sigma of the transverse beamsize, σ .

The RMS of the beam motion is different depending on the location around the ring and ranges between 0.6 μ m up to 3.5 μ m. To identify the frequencies that contribute to the overall motion, Figure 1 shows the integrated spectrum of the beam motion collected during regular operations from the storage ring BPM electronics, Instrumentation Technologies' Libera Electron (Libera), at a sampling rate of 10 kHz. The largest contributor to the beam motion is at the mains frequency at 50 Hz and its harmonics with mechanical vibrations contributing to the noise below 100 Hz. The recent improvements to the beamlines have meant that some are now seeing the effects of the 50 Hz beam motion.

To improve the beam stability a fast global orbit feedback system was proposed. Fortunately there are many existing solutions to learn from; CLS [1], Elettra [2], SLS [3], Soleil [4], SPEAR III [5], TLS [6] and NSLS II [7]. One requirement desired for the design was to have a single centralised processing station. The number of BPMs in the feedback (98 in total) was sufficiently small to make this possible. This has the potential to reduce hardware costs and simplify the design of the system where synchronisation of multiple processing stations is not needed.




Figure 1: Integrated spectrum of the beam motion at ID5 (IVU). The vertical beamsize is based on the model with the current emittance coupling value of 1%.

The proposed layout of feedback system is shown in Figure 2 where the outer fibre optic ring is used to aggregate the position data from all the BPMs and the inner ring is used to distribute the corrections to the corrector power supply units (PSUs). A 1000 Base-T GbE connection links the outer ring with the inner ring through the central processor.




Figure 2: Proposed layout of the feedback system forming two fibre optic rings (2 Gbps). The FPGAs send corrections to the corrector PSUs via serial lines. The PSU, Liberas and FPGAs have network connections to the control system.

The design of the system is broken into three parts that shall be discussed in the following sections: position aggregation, feedback processing and the corrector magnets.

POSITION AGGREGATION

With the Liberas there were two options for the aggregation of position measurements at 98 BPMs at a rate of 10 kHz. The first is the communication controller (CC) developed at DLS [8] and Libera grouping [9]. The choice was made to use the Libera grouping because the combination of the Gigabit Ethernet and UDP/IP protocols opens the possibility of choosing many different platforms for feedback processing. I-Tech was commissioned to extend the grouping size from 64 to 128 Liberas. Each Libera sends adds 128 bits to the data packet with a 336 bit header giving a total of 12880 bits in one UDP packet. By grouping all BPMs into a single data packet, much like the CC, the synchronisation and aggregation are handled transparently by the Liberas and simplifies the design of the central processor. The expected contribution to the total latency of the feedback system is approximately 67 μ s for 98 BPMs, including the packet encoding and transmission time. This system has been commissioned and has been operational for few months.

Two Liberas are currently used to stream the position data on a dedicated virtual local area network (VLAN) to the prototype feedback processor and to a PC in the control room respectively. The PC in the control room runs the fast acquisition archiver [10] extended to accept the UDP packets. The archiver is used to access the fast acquisition data and is used as a beam stability monitor.

FEEDBACK PROCESSING

An FPGA solution was chosen to do the feedback processing because it was the most suitable as a robust real-time platform to perform the calculations. Other real-time systems such as RTEMS and RTAI [11] were also considered. However the FPGA platform was favoured because of its processing potential and within the facility it was a platform that we wanted more experience in developing.

The position data grouped by the Liberas are sent to a prototype feedback processor via Gigabit Ethernet. The processor is an FPGA development board by Hitech Global (V6 PCIE) with a Xilinx Virtex-6. Here a state-machine hardware design in the FPGA parses each packet to extract status and position information from each BPM. A processor-less design in Verilog HDL (hardware description language) approach was taken because it offers the advantage of being able to run the calculations (and other functions) in parallel hardware without the limitations imposed by the inherently sequential operation of a microprocessor.

As each BPM's unique identifier, i , and X and Y beam position information is extracted from the Ethernet packet, it is plugged into a running sum type calculation,

for each unique corrector, n , using the following algorithm:

```
for(n=0;n<N_CORRECTORS;n=n+1) begin
    corrector[n] <= corrector[n] +
        coeff_X[n][i] * libera_x_pos[i] +
        coeff_Y[n][i] * libera_y_pos[i];
end
```

where $coeff_X$ and $coeff_Y$ are the coefficients of the inverted BPM-Corrector response matrix. It should be noted that the Libera data in the packet doesn't arrive in a nicely ordered sequence according to its unique ID, i .

The values for $coeff_X$ and $coeff_Y$ are selected out of block RAMs on the basis of the i . These RAM blocks have to be arranged, one per corrector, so that they can be accessed simultaneously for all of the correctors. This is because the for-loop is flattened during logic synthesis and the calculations for all the correctors run in parallel. This means that when the last set of X and Y positions for the last Libera become available, the matrix calculation is complete, with very little delay. Care must be taken to then check that the Ethernet packet was valid and not to send the calculations on if it wasn't.

The communications system chosen to distribute the calculated corrector values to the FPGA devices controlling the corrector magnets is the bi-directional Xilinx Aurora 8B10B protocol (the 8B10B corresponds to the line coding used) which is a high-speed serial connection that can be simply implemented into a FPGA design from a supplied HDL module using the Xilinx synthesis tools. It is simple to integrate, involving nothing much more than writing the values to be sent into a fixed length register and setting a transmit bit at the transmit end, and at the receive end, waiting for a receive bit to go true then reading the data out of a register. The physical media chosen is fibre-optic with a line rate of 2.5Gbps. The Virtex-6 FPGA we are using has multiple GTXs (fast serial transceivers) available to use and the development board (also Hitech Global V6 PCIEs) connects two of these GTXs to SFP (small form pluggable) adapters that can take a fibre-optic hardware module to supply the physical interface to the fibre-optic cable pairs used. We need two SFPs on each of the corrector magnet controllers so we can implement a daisy-chain network. Each corrector controller takes the relevant corrector values from the data arriving by the first SFP port and passing the rest of the corrector values down the line on the second SFP port. An addressing scheme will be implemented to identify which values correspond to which corrector.

CORRECTOR MAGNET AND PSU

The corrector magnets in the SR are trim coils on the sextupole magnets. These are powered by Danfysik 7000 90/120 Amp bipolar power supplies with a measured response of < 100 Hz. Any modification to the current PSU was deemed too much of a risk to operations. So the possibility of adding a trim power supply in parallel to the

main corrector PSU was also considered. This however would still involve modifications to the main supply to disable the feedback loops to let the trim supply “take over”. This was also considered unacceptable. Therefore the decision was made to separate the fast and slow PSUs and add secondary corrector coils on the sextupoles for the fast correctors. To ensure that this can be done a number of tests were carried out.

Vacuum Chamber Bandwidth

At the ASLS the vacuum chambers are made from 3 mm thick stainless steel. The effective attenuation of the fields from the stainless steel chamber was calculated using the skin depth equation given by

$$\delta = \sqrt{\frac{\rho}{\pi f \mu_r \mu_0}} \quad (1)$$

where δ is the skin depth, ρ the resistivity, f the frequency, μ_r the relative permeability and μ_0 the permeability of free space. At one skin depth the metal’s conductivity is reduced to 36.8% and at two skin depths its 13.5%. To calculate the transmission of the AC magnetic field penetrating the vacuum chamber the following equation is used,

$$A = \exp(L / \delta) \quad (2)$$

where A is the fraction of the magnetic field transmitted and L the vacuum chamber thickness. For a 3 mm thick stainless steel plate. It is anticipated that the field will attenuate by 71% at ~2 kHz for an infinite plate model.

To measure the response of the vacuum chamber a Danfisk PSUs was replaced with 100W bipolar Kepco power supply with a maximum current of 5 A. The Kepco was controlled by an external voltage from a frequency generator. The Kepco’s output was configured to supply a sinusoidal current with a peak of 230 mA with frequencies ranging from 1 Hz to 1600 Hz. The perturbation to the beam was measured by analysing 10 kHz position data from a single BPM in the SR. The amplitude of the frequency peak at the drive frequency is plotted in Figure 3. The results show that the oscillation amplitude, and indirectly the magnetic field, decreases by 71% at 350 Hz (horizontal) and 1100 Hz (vertical). The bandwidth is wide enough given the frequencies of the perturbations are < 500 Hz.

Magnet Inductance

During the measurement of the response of the vacuum chamber the peak voltage on the Kepco was also measured. With increasing frequency there was a linear increase in the peak voltage due to the inductance of the magnet. Using $V = L\omega$ and the measured data, the inductance was calculated to be 1.7 ± 0.1 mH (horizontal) and 0.8 ± 0.1 mH (vertical). This compares reasonably well with theoretical calculations of 1.96 mH (horizontal) and

0.65 mH (vertical). The low inductance value reduces the voltage, and power, requirement of the power supply.




Figure 3: Amplitude of the peak at the drive frequency when a corrector magnet is driven by a sinusoidal current with a peak current of 230 mA. The vertical magnetic fields that drive the horizontal motion attenuates faster than the horizontal magnetic field due to the rectangular geometry of the vacuum chamber. The multiple lines are from different data sets. The variation between data sets may be due to the instability of the Kepco driving an inductive load.

Magnet Strength Requirement

The largest perturbation is measured at 50 Hz and the required correction is $2.6 \mu\text{rad}$. For perturbations induced by insertion device (ID) motion there are feedforward tables for the local ID correctors. However during large movements there will be uncorrected perturbations. These are expected to be $< 1 \mu\text{rad}$ because the largest rate of change among all the local ID correctors is $6 \mu\text{rad/s}$. As for slow drifts, over a 2 hour period, the maximum change in the correctors (slow orbit feedback currently being used) is at most $4 \mu\text{rad}$ (in both decay and top-up operation). Therefore the fast corrector will be designed to provide a maximum correction of $10 \mu\text{rad}$. For a 10 turn secondary corrector coil around the sextupole this would require a current of 0.8 A.

Fast Corrector Trim Coils on Sextupoles

There was concern that the slow and fast PSU would couple given transformer like arrangement on the sextupole. A test with a custom built 1A low power amplifier indeed showed a significant coupling between the slow and fast PSUs. The solution is to ensure that the magnetic field created by the slow correctors are perpendicular to the fields created by the fast correctors. In this setup the coupling between the PSUs was not detected. In the SR there are a total of three slow

horizontal and four slow vertical correctors. Using additional coils there can be 4 fast horizontal and 3 fast vertical correctors.

To decrease the installation time a multi-core cable with connectors on either end will be designed such that when the two connectors are joined the inner wires of the multi-core cable will form a multi-turn coil.

Fast Corrector PSU

Given that the requirement for the corrector is only 0.8 A and low voltage requirement because of the low inductances, the power supply would be quite straight forward to design. A PSU was built to test the coils and the exercise was useful in determining the cost and effort required to build all 98 supplies including the interfaces. There are also commercial sources such Bilt system [12] that is used by Soleil and ESRF for their feedback system that is perfectly suited for this purpose.

CONCLUSION

Thus the feedback system will consist of:

- Libera grouping providing synchronised position data at a rate of 10 kHz via GbE.
- Position data is processed in an FPGA (currently only a matrix multiply algorithm) ready to be distributed to “slave” FPGAs connected to power supplies.
- The power supplies will need to be able to drive a minimum of 1 A into a magnet with an inductance of ~2 mH.
- The magnet will be secondary corrector coils wound around the sextupole magnets with a minimum of 10 turns.

One of the defining parameters for the bandwidth of the system is the total latency of the system. This is expected to be very similar to Soleil’s calculations [4] which is estimated to be 360 μ s.

REFERENCES

- [1] M. G. Abbott et al., EPAC, THPC118, pp3257-3259, 2008. M. T. Heron et al., ICALEPCS, THB003, pp650-652, 2009.
- [2] M. Lonza et al., PAC, MOPAN024, pp203-205, 2007.
- [3] V. Schlott et al., Proceedings of BIW, 2004.
- [4] N. Hubert et al., DIPAC, POT019, pp187-189, 2005. N. Hubert et al., DIPAC, TUPC20, pp189-191, 2007. N. Hubert et al., BIW, WECOTC02, pp337-341, 2008. N. Hubert et. al, DIPAC, MOOC01, pp27-31, 2009.
- [5] A. Terebilo and T. Straumann, EPAC, THPCH102, pp3035-3037, 2006.
- [6] C. H. Kuo et al., DIPAC, WEPC17, pp346-348, 2007.
- [7] Y. Tian and L. H. Yu, PAC, WEODN4, pp1-3, 2011.
- [8] I. S. Uzun, CTRL-FOFB-CC-0001, 2009.
- [9] A. Bardorfer and T. Karcnik, EPAC, pp1041-1043, 2008.
- [10] M. G. Abbott et al., ICALEPCS, THCHMUST03, pp1244-1247, 2011.
- [11] E. Bianchi and L. Dozio, “Some experience in fast hard real-time control in user space with RTAI-LXRT” Real Time Linux Workshop, 2000.
- [12] http://www.bilt-system.com/appli_acc.html (25/09/2012)