A NEW DIFFERENTIAL AND ERRANT BEAM CURRENT MONITOR FOR THE SNS* ACCELERATOR

W. Blokland, C. Peters, ORNL, Oak Ridge, TN 37831, USA

Abstract

A new Differential and errant Beam Current Monitor (DBCM) is being implemented for both the Spallation Neutron Source's Medium Energy Beam Transport (MEBT) and the Super Conducting Linac (SCL) accelerator sections. These new current monitors will abort the beam when the difference between two toroidal pickups exceeds a threshold. The MEBT DBCM will protect the MEBT chopper target, while the SCL DBCM will abort beam to minimize fast beam losses in the SCL cavities. The new DBCMs will also record instances of errant beam, such as beam dropouts, to assist in further optimization of the SNS Accelerator. A software Errant Beam Monitor was implemented on the regular BCM hardware to study errant beam pulses. The new system will take over this functionality and will also be able to abort beam on pulse-to-pulse variations. Because the system is based on the FlexRIO hardware and programmed in LabVIEW FPGA, it will be able to abort beam in about 5 us. This paper describes the development, implementation, and initial test results of the DBCM, as well as errant beam examples.

INTRODUCTION

Errant Beam

Errant beam at SNS is any beam outside the normal operation envelope. For example, beam with a too high peak density that can damage the spallation target, any beam that causes high losses, or too much beam scraped into the MEBT chopper target. In some cases, we don't know if certain anomalies, like a few µs of Low Energy Beam Transport (LEBT) chopper misfires, are damaging, but we want to investigate and document these anomalies anyway to further our understanding of the accelerator.

Impact on SCL Performance

Examples of errant beam in the SCL are abrupt beam losses caused by low current or truncated beam pulses, as well as beam pulses with incorrect energy. Beam losses cause SCL cavity damage by releasing contaminants or condensed gasses that then absorb RF power and become ionized. Discharge or arcing can then occur which can damage the niobium/copper/ceramic surface of the cavity. Such surface damage is accumulative and to prevent further damage the RF field must be lowered leading to lower overall energy for the accelerator, see [1]. Even short beam losses, less than 15µs are now thought to contribute to this damage.

*ORNL/SNS is managed by UT-Battelle, LLC, for the U.S. Department of Energy under contract DE-AC05-00OR22725

Initial Capturing of Errant Beam Events

An investigation was done to measure how much beam was being lost, how often it was occurring, and what the causes were. The goal was to reduce the frequency of the errant beam, and then to reduce the impact of any remaining errant beam.

Data was acquired from current monitors, RF waveforms, and loss monitors. BCMs were modified to acquire at the full beam rep rate and to be able to compare pulse-to-pulse variations. If there was a latched beam abort and the beam stayed off for >2 sec, then the BCMs pushed data to a file server while a console script stored RF and BLM waveforms. A tool, see Fig. 1, was developed to correlate and view the data, see also [2].

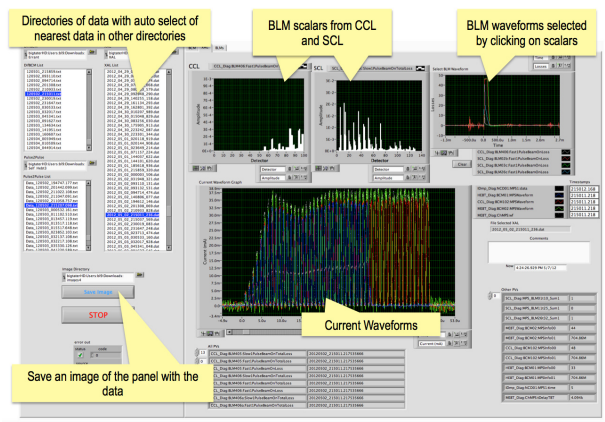

Figure 1: Errant beam data viewer showing a typical abort on SCL beam loss.

The data showed that the errant beam events had losses for about 15-20 µs until beam was shutoff by the Machine Protection System (MPS), and that these events were occurring about 30 times per day. The causes for the events were the Ion Source/LEBT and the Warm Linac RF cavities, with the RF being the main contributor (>90% of the faults).

Because issues with Warm Linac RF caused the majority of the faults, our focus was to reduce the frequency of these trips. Adjustments were made to the RF (gradient changes, resonant frequency changes, and preventative maintenance on vacuum systems), which resulted in the fault frequency being reduced by more than a factor of two. SCL downtime was reduced by a factor of six, but because the Warm Linac RF faults cannot be completely eliminated and the damage to SCL cavities due to the errant beam continues the final goal is to reduce the impact of the errant beam in the SCL, see [3].

ISBN 978-3-95450-127-4

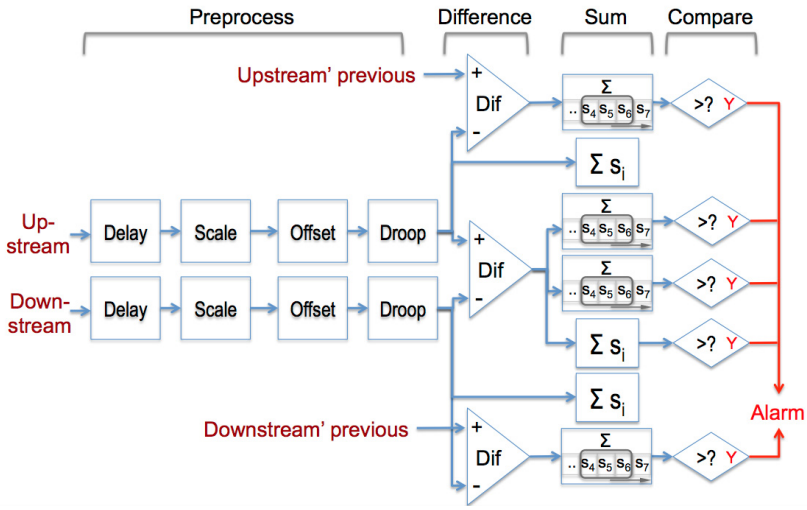

Figure 2: Diagram of the FPGA functionality.

IMPLEMENTATION

The BLM (Beam Loss Monitor) and MPS (Machine Protect System) have been designed to abort beam in about 20 µs. While the BLMs and the MPS could be upgraded, a single DBCM (Differential Beam Current Monitor) can protect the whole SCL and also detect errant beam other than beam losses, such as missing pulses or truncated beam. The requirements for the DBCM are to abort beam as fast as possible but because of the distances involved, the main goal is set to abort the beam in about 5 µs when all beam is lost.

The existing MEBT DBCM, which protects the chopper targets and scrapers, is based on obsolete hardware and our group no longer has the resources to do VHDL language FPGA programming and adapt the program to the new requirements.

The FlexRIO-based hardware from National Instruments allows us to program the FPGA in LabVIEW style programming, which is well supported in our group. For the new DBCM, we choose a PXIe-based system with a LabVIEW RT (Real-Time) controller communicating with the FPGA mounted with a 14-bit 100MS/s ADC.

FPGA Processing

To compare the two beam currents before and after the SCL, we picked the last current transformer before and the first one after the SCL: the CCL102 and the HEBT01 current transformers. The transformers have a 1 GHz bandwidth but also have an approximate 1 ms droop time constant and thus show significant droop during the 1 ms long beam pulse. The FPGA processing hence includes a droop correction filter in addition to the delay correction to adjust for cable length differences, a scaling function to correct for differences in attenuation or polarity, and an automatic offset correction to compensate for low frequency drift. Each preprocessed sample is stored, sample-by-sample, in local buffers, to be recalled for the next beam pulse to do the pulse-to-pulse comparison. Next, three differences are calculated, between the upstream and downstream current and between the current and previous current pulse for both the upstream and downstream currents. Sliding windows create sums of these differences. The FPGA also calculates the total sum of each channel and the channels' difference. The results from the sliding windows and the sum of the upstream and downstream difference are compared to a threshold to decide whether to alarm or not, see Fig. 2.

Up to three selected waveforms and the alarm status waveform are sent to the RT controller during the acquisition. At the start and end of each waveform is a header and a footer with information about the number of samples, sampling frequency, trigger event, cycle ID, EPICS timestamp, and total integrated charge of the upstream and downstream current waveforms.

The trigger event, cycle ID, and EPICS timestamp are only available if the SNS Timing (EVNT) and data (RTDL) links are hooked up to the FlexRIO board. SNS does not have PXI-based timing decoder board. However, we found that we can directly implement the decoding functions on the FlexRIO FPGA [4]. Figure 3 gives an example of preprocessing and SNS Data link receiver LabVIEW FPGA code. The programming style is similar of the regular LabVIEW data flow but different in that all processing is pipelined and must be optimized to reach clock rates of 100 MHz.

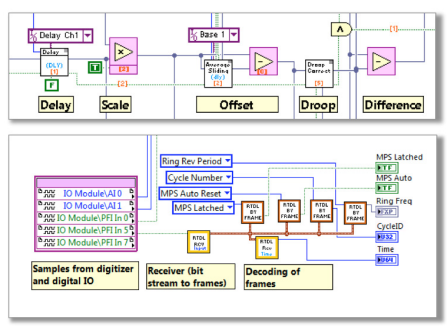

Figure 3: The preprocessing (top) and RTDL receiver (bottom) LabVIEW FPGA code.

ISBN 978-3-95450-127-4

Real-time Processing

As all of the time-critical features have been implemented in the FPGA, the Real-Time system only has to calculate statistics and present the data to the control system. We used a native LabVIEW Channel Access library to implement the required Process Variables (PVs), to communicate with the EPICS-based control system, see [5].

Most PVs are not explicitly declared in the RT code, but rather are derived from the LabVIEW data structures used in the program. This way, if the LabVIEW program is changed, the new PVs are automatically declared and instantiated, saving the programmer the work of doing this explicitly. Similarly, the program automatically creates the appropriate Control System Studio (CSS) OPerator Interface (OPI) files to function as initial control room console screens. For the final screens, the programmer must still rearrange the OPI display elements using the OPI editor. The screens display the selected waveforms and allow the user to set up the FPGA processing and RT statistics, see Fig. 4.

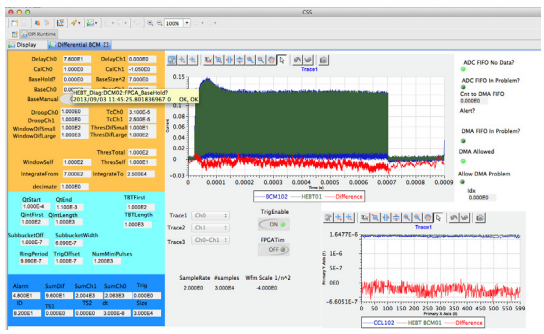

Figure 4: OPI screen of the differential BCM.

The RT program receives up to three waveforms from the FPGA to display. The user can select the main waveforms, the upstream, downstream, and the difference of the two, but also intermediate values from the preprocessing or values calculated by the sliding windows, to help with the debugging and the setting of the thresholds for the alarms. The RT program currently stores the last 30 waveforms that were tagged with an alarm and but only the last tagged waveform is displayed through the OPI displays.

INSTALLATION

Since we have already an older style DBCM in the MEBT and minimizing the damage to the SCL is the most important issue, we placed the new DBCM in the SCL. Cables were pulled to bring the signals from the High Energy Beam Transfer (HEBT) BCM01 and CCL BCM102 to a location in the SCL service building. To make up for the additional cable length and extra noise induced on the cables, we amplified the signals in the electronics buildings and then attenuated at the DBCM, see Fig. 5.

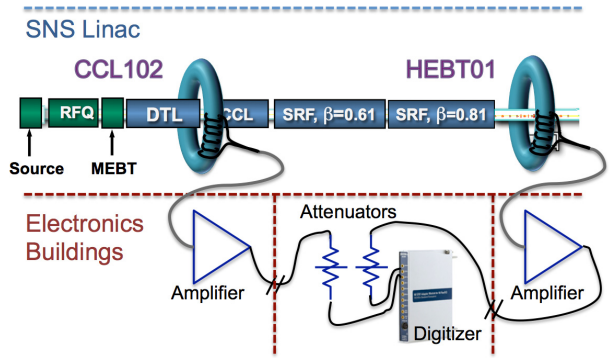

Figure 5: Cabling layout of the differential BCM.

INITIAL RESULTS

While the Machine Protect System (MPS) has not been hooked up yet to the DBCM, we can already see when the system would have aborted the beam. For each sample of the beam current waveform, the system also provides the alarm status. Figure 6 shows an example of the beam loss condition we are looking to abort more quickly to avoid damage to the SCL cavities. Beam goes into the CCL but does not arrive at the HEBT; it is either lost in the SCL or in the CCL, in either case possibly damaging the SCL cavities. The figure shows that first the alert on the small window difference occurred followed by the downstream pulse-to-pulse difference (the previous pulse was full length). Finally, the upstream pulse-to-pulse alert was given because the beam had now been aborted.

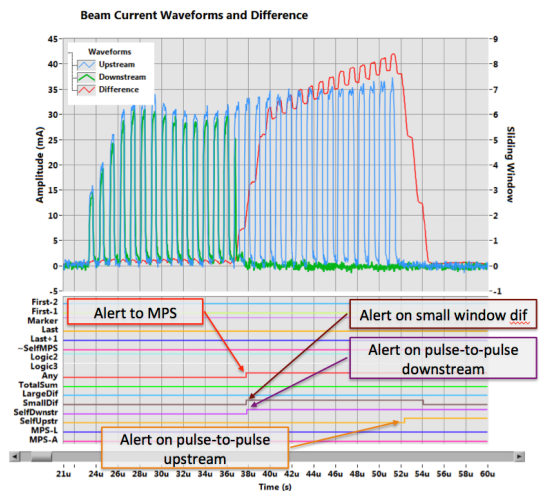

Figure 6: An example abort in the SCL due to beam loss. The waveforms displayed are the upstream and downstream beam current waveforms and the short sliding window waveform of the difference. The MPS alert is the OR'd combination of all other alerts.

Figure 7 shows the DBCM detecting a truncated beam; this is probably due to a drop in the RFQ RF field. The system can also detect small temporary dips in the beam,

ISBN 978-3-95450-127-4

as those differ from the previous beam pulse, see Fig. 8. Figure 9 shows an instant dip in two pulses. This is due to misfiring of the LEBT chopper and does not cause losses in the SCL directly, but can lead to the RF feed-forward loop learning the wrong beam conditions and could lead to losses on the next pulse. If the drop lasts longer then a few µs then the RF field increases due to the change in beam loading, resulting in different beam energies, which can then lead to downstream losses. This is one of the conditions we hope to learn more about with the DBCM.

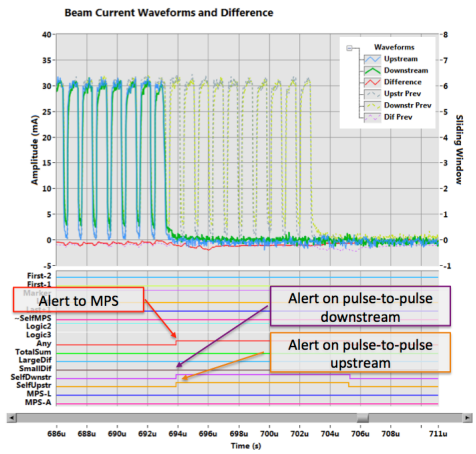

Figure 7: A truncation of the beam detected by the pulse-to-pulse difference without losses in the SCL section.

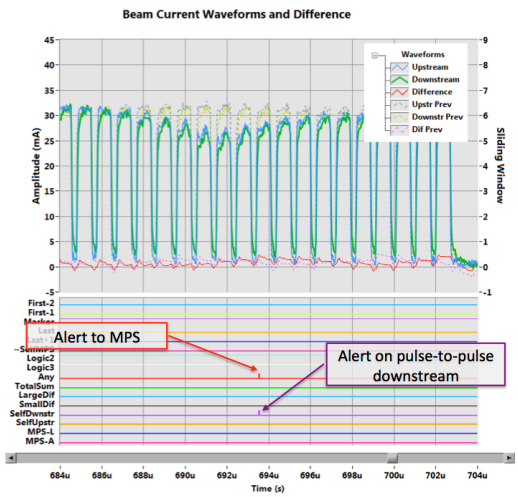

Figure 8: A temporary dip in the beam current detected by the pulse-to-pulse difference.

With the used threshold settings, the DBCM can abort in about 1 µs. With the existing long cabling, we estimate about 2.5 µs of signal cable delays and 1 µs of abort cable delays. Adding in about 1.5 µs beam propagation time gives a total delay of about 6 µs.

DISCUSSION

The LabVIEW FPGA based timing decoding library gives us a very flexible way to implement the timing decoding without having to develop and maintain a separate custom timing board. The library also works on the cRIO platform. We expect that with all the group's programmers able to develop on this platform that we have a long-term solution for many high performance diagnostics.

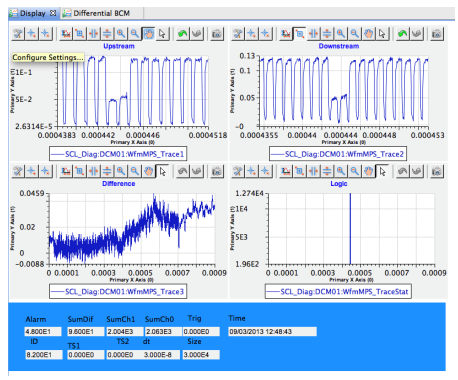

Figure 9: An immediate dip in two minipulses displayed through the CSS interface.

In the near future, we plan to hook up the DBCM directly to the front-end MPS input and, given the initial results, we expect to be able to abort with in about 6 µs. In the long term, we plan to install new transformers and pull cables from the tunnel directly to the DBCM location to reduce the signal cable delays and improve signal-to-noise allowing lower thresholds and thus abort in *about* 5 µs. As data is acquired and analyzed offline, we expect to add more statistics functionality to the real-time processing to get a better idea of the accelerator performance over time and to diagnose intermittent issues, such as a single half pulse or a small dip in the beam current.

REFERENCES

[1] S. Kim, et al., "The Status of the Superconducting Linac and SRF Activities at the SNS," 16th International Conference on RF Superconductivity, Sep 23-27, Paris, (2013)
[2] W. Blokland, "Errant Beam: Tools and Data," SNS Errant Beam Committee, May 2012, Oak Ridge, TN.
[3] C. S. Peters, "Errant Beam Update," Accelerator Advisory Committee, May 7, 2013, Oak Ridge, TN.
[4] R. Dickson, "LabVIEW FPGA SNS Timing User Guide," SNS, May 7 2013, Oak Ridge, TN.
[5] S. Zhukov, "Pure LabVIEW Implementation of EPICS Communication Protocol," Big Physics and Science Summit at NIWeek, August 7-8, 2012. Austin, TX.

ISBN 978-3-95450-127-4