

MACHINE LEARNING FOR TIME SERIES PREDICTION OF AN ACCELERATOR BEAM TO RECOGNIZE EQUIPMENT MALFUNCTION*

C. C. Peters[†], W. Blokland, D. Brown, F. Liu, C. Long, D. Lu, P. Ramuhalli, D. Womble, J. Zhang, A. Zhukov, Oak Ridge National Laboratory, Oak Ridge, TN, 37831, USA

Abstract

The Spallation Neutron Source (SNS) is an accelerator based pulsed neutron source based on a 1 GeV pulsed proton Superconducting Radio Frequency (SRF) linear accelerator (linac). Since beginning high power beam operation in 2006 correlations have been found linking abrupt beam loss events to SRF cavity instabilities. With the planned upgrades to double the beam power we expect increased rates of degradation and the importance of minimizing these beam loss events will become ever more important. To further limit degradation, we are developing machine learning approaches to monitor the beam and to detect, predict and prevent beam loss events. Initial research has shown that precursors to beam loss events are detectable. The initial steps are to use Machine Learning (ML) based classification to recognize anomalies and to use Long Short-Term Memory (LSTM) autoencoders to predict beam loss. In this paper, we describe recent progress in applying machine learning for recognizing anomalies and predicting beam loss and present initial results of our research using acquired data from different diagnostics and the Machine Protection System (MPS).

INTRODUCTION

The Spallation Neutron Source (SNS) Superconducting Radio Frequency (SRF) linear accelerator (linac) provides the highest power pulsed proton beam in the world. The linac routinely provides beam power of 1.44 MW for neutron production.

SRF Cavity Degradation

SRF cavity degradation means that over time, to maintain high availability, an SRF cavity gradient must be continually reduced. There are many causes for having to lower an SRF cavity gradient, but the one that has had the most focus has been beam loss. The beam loss can be abrupt like from an RF cavity fault or a slower process like small amounts of loss from beam halo. These types of events have been labeled errant beam [1].

* This manuscript has been authored by UT-Battelle, LLC under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

[†] peterscc@ornl.gov

The beam power ramp up for the SNS began in 2006 (see Fig. 1). The power was increased slowly reaching about 1 MW in late 2009. At that time some of the SRF cavities began to experience reliability issues related to beam loss. This initial instance of SRF cavity degradation was found to be caused by malfunctioning within the Machine Protection System (MPS) [2, 3].




Figure 1: Beam power ramp up history at SNS.

The issues with the MPS were resolved but the SRF cavity degradation continued, though at a reduced rate. Further studies revealed that the continuation of degradation was related to beam loss caused by faults from the Drift Tube Linac (DTL) and Coupled Cavity Linac (CCL) RF cavities. This area of the linac is labeled the warm linac due to the DTL and CCL RF cavities operating at room temperature relative to the SRF cavities, labeled the cold linac, operating at 2 Kelvin [4]. During a warm linac cavity fault beam loss occurs within a microsecond. The MPS beam turn off time is about 15 microseconds so during a warm linac RF cavity fault the beam will stay on and all of the beam will be lost for 15 microseconds. One thing to keep in mind is that there is no system to limit the peak beam current. If the peak current is increased the amount of beam lost during an event will increase by that amount. For beam power upgrades this may become an issue.

In addition to faults from the warm linac it was found that malformed beam pulses originating in the ion source also contributed to SRF cavity degradation. The ion source can trigger issues in multiple ways. The Low-Level Radio Frequency (LLRF) system controls the cavity field amplitude and phase. The system has both feedback and Adaptive Feed Forward (AFF) [5]. Feedback controls the RF amplitude and phase during the pulse. When the ion source malfunctions during the pulse the extra RF power supplied to the cavity for beam acceleration is converted

into cavity RF gradient until the feedback can lower the supplied RF power. The increase in gradient increases field emission exponentially and can trigger a cavity to start quenching if the field setting is near instability. The more the ion source is malfunctioning the more times the field will be pushed into instability. The AFF measures RF field amplitude and phase errors pulse to pulse at a 20 Hz rate (every third pulse). The system then applies the error correction for the next 3 pulses. The AFF system will not learn the errors for application if an MPS fault occurs on the learning pulse. This is an important fact because the ion source equipment is not connected to the MPS. This means the ion source can malfunction and the AFF will learn the errors on a malformed beam pulse and apply the correction to the following pulses. This results in significant beam loss.

Lastly there is beam halo caused by improperly tuned beam (ion source high voltage or linac magnets). The mistuning introduces a small increase in beam loss slowly over time that can increase heating within SRF cavities and again push cavities into instability. The degradation for one SRF cavity is shown in Fig. 2.




Figure 2: SRF cavity 06a field has been lowered from 12.5 MV/m to 10.2 MV/m over the last 7 years. The field has stabilized since installing a new interlocking scheme in 2018.

LIMITING THE IMPACT OF ERRANT BEAM

After the MPS issue was resolved and SRF cavities continued to degrade even with the nominal 15 microsecond beam turn off time a new system was developed to turn the beam off faster.

The new system, the Superconducting Cavity Linac (SCL) Differential Current Monitor (DCM) [6], uses a Beam Position Monitor (BPM) in the first part of the CCL and a Beam Current Monitor (BCM) at the end of the SCL. The system measures the beam current at each location during the beam pulse and subtracts the two and compares to an adjustable limit. When the difference exceeds the limit, it sends a signal to interrupt beam. The response time of the system is on the order of 1 microsecond. The system connects directly to the Low Energy Beam Transport

(LEBT) chopper and begins chopping the beam before it enters the RFQ. While the beam is being chopped the LEBT chopper control sends a signal to the MPS to turn off the beam. The direct connection to the LEBT chopper cut the beam turn off time in half to about 8 microseconds.

Even with the reduction in turn off time some of the SRF cavities still required gradient reductions to maintain high availability. This was because nothing was in place to trigger the MPS to fault when an ion source malfunction occurred.

The SCL DCM system software was upgraded to compare beam waveforms pulse to pulse at each location (CCL and SCL). The logic is, if the new pulse has less current than the previous the system will fault through the same mechanism as the upstream-downstream difference mentioned previously. This solved the issue of the AFF learning on malformed ion source pulses. The change was implemented in 2018, and since that time the gradient for SRF cavity 06a has not needed to be decreased since implementation.

MACHINE LEARNING TO PREVENT ERRANT BEAM LOSS

The SCL DCM solution has worked very well but looking forward to the Proton Power Upgrade (PPU) [7] there will likely be a need to either reduce the turn off time further (will be very difficult) or develop a new method to limit beam loss. Instead of focusing on a new design to shorten the MPS turn off time the focus recently has been on trying to use machine learning to try to predict when an errant beam pulse is going to occur and just do not send beam for that pulse/pulses.

Machine Learning Operational Requirements

The accelerator at SNS is pulsed 60 times per second. This means there are a possible 5,184,000 beam pulses. There are many different beam trip lengths, but they can be binned to determine averages. The beam trip rates that machine learning will focus on will be the trips that last less than 1 minute. These are the errant beam pulses that require operator intervention to reset. Current less than 1 minute trip rates are ≈ 13.6 trips per day (see Fig. 3).




Figure 3: Binned beam trip frequencies since FY09. Reducing the first bin is the goal for machine learning.

On average these trips last ≈ 41.3 seconds. On an average day this means $\approx 33,706$ pulses are lost due to errant beam faults. This is $\approx 0.65\%$ of the possible daily beam pulses lost. From a beam on target loss of neutrons perspective

machine learning needs to achieve detection close to this percentage including both true negative (bad pulses labelled bad) and false negative (good pulses labelled bad) to make it operationally feasible.

It is well understood that beam loss events do degrade SRF performance, but it is difficult to estimate how much of a reduction in downtime will occur by limiting the number of errant beam pulses.

The current SCL DCM system can detect beam abnormalities and turn the beam off approximately 8 microseconds after beam loss begins. To improve on this, machine learning only needs to detect an anomaly approximately 10 microseconds before the beam loss begins to occur. This will leave enough time for signal travel time to the beginning of the accelerator to disable beam. Nominally being able to detect the bad pulse an entire beam pulse ahead of time should solve all the issues seen from errant beam, and that has been the focus thus far.

Initial Machine Learning Results

When the SCL DCM was installed, the initial setup already had the ability to store the beam current waveform for the pulse before the errant beam pulse as well as the pulse after the errant beam pulse (see Fig. 4). Data are from both the BPM in the CCL as well as the BCM at the end of the SCL.




Figure 4: Example of the SCL DCM data storage from the initial installation. The first set of waveforms are before errant beam, middle during, and on the right after the event. All show beam current during the beam pulse.

Waveform data from 2015 were analysed to determine whether there are signatures in the bad pulses that can predict errant beam on the following pulse [8]. Multiple algorithms were used (k-nearest neighbor, decision tree classifiers and regressors, random forests and gradient boosts, linear support vector machines, perceptrons and neural networks, and others). The best performing classifier on a general dataset was the logistic regression analysis with 79.5% test precision for true negative prediction. When data were filtered based on metadata and analysed using neural network the true negative score did reach up to 91.9%. The issue is the number of false negatives were at 8% which is an order of magnitude too high.

The important result of the initial analysis is there are markers in the beam pulse before the errant beam event that can be used to predict the expected health of the next pulse.

Further Analysis and Future Plans

The initial analysis was done using data from 2015 so it was not clear whether the same analysis would produce the

same results. The analysis was redone on data from July 2020. The data produced the same results at about 80% prediction rate for true negative. False negative results were not analysed.

The initial analysis also presented the need to filter beam waveform data based on the accelerator state. Events can now be filtered based on the equipment that caused the fault. Waveform data for a particular piece of equipment should all have the same fault signature and may improve prediction accuracy.

BPM phase data are also stored but no machine learning analysis has been done on the BPM data. Figure 5 shows an example of BPM data where pulse to pulse phase data instabilities are clearly visible by eye. These data should provide an easier starting point for machine learning experts to understand accelerator data and look for signatures.




Figure 5: Example of the BPM phase data during the beam pulse. The top is just upstream of a set of SRF cavities and the bottom just downstream of the same set of SRF cavities. The different colors are 180 different beam pulses. Downstream of the cavities the phase slews enough to be seen by eye. This event caused a beam trip.

To that point the data filtering and capture are done by accelerator experts. The machine learning analysis is done by machine learning experts. Progress has been serial. Machine learning experts waiting on data and then accelerator experts waiting on analysis. The next focus is for accelerator experts to begin performing the machine learning analysis as well.

CONCLUSION

Significant improvements have been made to reduce both the frequency and impact of errant beam in the SCL at SNS. Initial results of using machine learning for bad pulse detection shows that the pulses can be predicted at the 80-90% level. The issue remains that there are too many false negative (labelling good pulses as bad) predictions.

REFERENCES

- [1] C. Peters *et al.*, "Understanding the Source and Impact of Errant Beam Loss in the Spallation Neutron Source Superconducting Linac", in *Proc. 61st ICFA ABDW on High-Intensity and High-Brightness Hadron Beams*, Daejeon, Korea, May 2018, paper MOP1WB02, pp. 48-53.

- [2] C. Sibley III *et al.*, “The SNS Machine Protection System: Early Commissioning Results and Future Plans”, in *Proc. of PAC’05*, Knoxville, TN, May 2005, paper RPPE021, pp. 1727-1729.
- [3] D. Curry *et al.*, “MPS Delay Measurements”, submitted for publication.
- [4] S. Henderson *et al.*, “The Spallation Neutron Source Accelerator System Design”, *Nucl. Inst. Meth. A*, vol. 763, pp. 610-673, 2014. doi:10.1016/j.nima.2014.03.067
- [5] K.-U. Kasemir, M. Champion, M. T. Crofford, and H. Ma, “Adaptive Feed Forward Beam Loading Compensation Experience at the Spallation Neutron Source Linac”, in *Proc. of PAC’05*, Knoxville, TN, USA, May 2005, paper MPPP016, pp. 1467-1469.
- [6] W. Blokland and C. C. Peters, “A New Differential and Errant Beam Current Monitor for the SNS Accelerator”, in *Proc. of IBIC2103*, Oxford, UK, Sep. 2013, paper THAL2, pp. 921-924.
- [7] J. Galambos *et al.*, “Status of the SNS Proton Power Upgrade Project”, in *Proc. 29th Linear Accelerator Conf. (LINAC2018)*, Beijing, China, Sep. 2018, pp. 24-28. doi:10.18429/JACoW-LINAC2018-M01P01
- [8] M. Rescic *et al.*, “Predicting Particle Accelerator Failures Using Binary Classifiers”, *Nucl. Inst. Meth. A*, vol. 955, p. 163240, 2020. doi:10.1016/j.nima.2019.163240