ISBN: 978-3-95450-214-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2021-THPAB252

MACHINE LEARNING FOR IMPROVED AVAILABILITY OF THE SNS KLYSTRON HIGH VOLTAGE CONVERTER MODULATORS*

C. Pappas†, D. Lu, Oak Ridge National Laboratory, Oak Ridge, TN, USA
M. Schram, Thomas Jefferson National Accelerator Facility, Newport News, VA, USA
D. Vrabie, Pacific Northwest National Laboratory, Richland, WA, USA

Abstract

The targeted beam availability at Department of Energy's Spallation Neutron Source (SNS) user facility is greater than 95%. The High Voltage Converter Modulators (HVCMs) that are used to power the linac klystrons have been one source of lost beam time. In this work, we explore the use of machine learning (ML) methods to address two problems in HVCMs to improve their reliability. First, prediction of aging and failure of HVCM components such as capacitors, rectifier assemblies containing hundreds of diodes, and insulating oil, in order to preemptively schedule maintenance and replacement. Second, automated tuning of the HVCM operation after service. The ML methodology pursued includes feature extraction, anomaly detection, and autoencoder network learning. The preliminary result indicates that ML can provide advance notice of upcoming faults. Our future work includes the improvement of anomaly prediction rate and HVCM component health state estimation and HVCM autonomous re-tuning.

INTRODUCTION

The HVCMs deliver cathode power to the klystrons which provide RF to the accelerator cavities. They convert 13.8 kV line power to ±1200 VDC, that is then feed to three H-bridge converters operating in parallel to produce up to approximately 135 kV, 1.3 ms pulses at 60 Hz [1]. A simplified schematic of a modulator is shown in Fig. 1. Prior to substantial upgrades to the HVCMs, they were the source of significant downtime as shown in Fig. 2 [2].

The HVCMs at SNS continue to occasionally experience catastrophic failures which can result in a day or more of lost beam time. Recent failures since 2018 that resulted in significant down time were due to failure of the resonant capacitors Ca, Cb or Cc, failure of the rectifier stacks Da1-Dc2 and overheating of components in the high voltage tank due to polymerization of the FR3 insulating oil. SNS uses a total of 14 HVCMs powering 92 klystrons as shown in Fig. 3.

The Proton Power Upgrade [3], will add an additional 3 HVCMs and 28 klystrons. Each HVCM uses a National Instruments PXI based controller to monitor various subsystems and communicate with the control room via ethernet. The controller has a total of 63 analog and 108 Boolean PVs, plus an additional 32 waveform channels. The waveforms are digitized at 50 MS/s, but decimated to 2.5 MS/s with a record length of 36 ms centered on the next to last output pulse. The waveforms representative to the last shot before shutdown are archived, as well as those characterizing the initial start up after a HVCM has had maintenance. Full bandwidth waveforms are saved with a record length of 3 ms, centered on the last pulse, whenever a fault occurs.

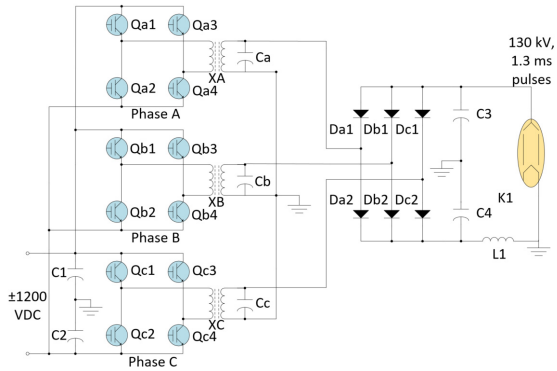

Figure 1: Simplified schematic of a HVCM.

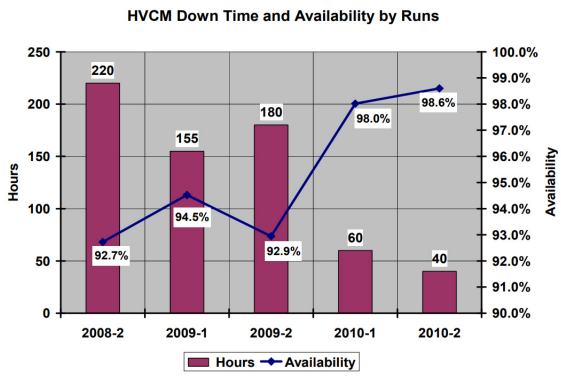

Figure 2: Early operating history of HVCM systems at SNS.

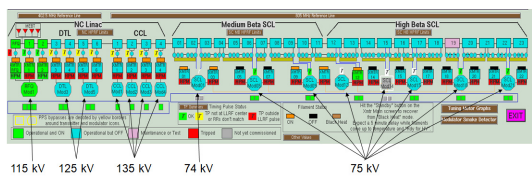

Figure 3: Layout of the SNS linac HVCMs and RF systems.

* This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 and DE-AC05-76RL0-1830 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE 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).
† pappasGC@ornl.gov

ISBN: 978-3-95450-214-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2021-THPAB252

In addition to the data available through the HVCM controller, SPICE circuit models have been built and verified to aide in design and troubleshooting. Using the models, it is possible to quickly simulate many types of fault conditions, or change the HVCM tuning without risk to equipment or operations. The HVCMs thus offer a rich source of data for machine learning (ML). An example of a simulation of the degradation of a resonant capacitor is shown in Fig. 4. In the simulation, the capacitor starts off with the nominal value until 20 ms, after which the value of the capacitor drops linearly to 50% nominal at 180 ms, then to 10% nominal after 190 ms. The repetition frequency has been increased to 200 Hz to provide greater granularity, and an ideal voltage source was used to prevent dropping of the capacitor bank. The impact of the change in capacitor is clearly seen on the IGBT current for the phase with the failing capacitor as shown. Other signals, not presented, are affected as well, including the transformer flux and the klystron cathode voltage.

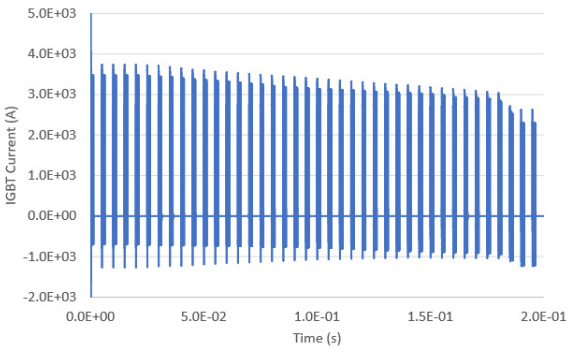

Figure 4: Results from SPICE simulation of failing capacitor.

The SPICE data collected for the various values of resonant capacitance can be utilized to train ML models in a supervised approach. These models can then be transferred and retrained on real data from the HVCMs and then used to observe degradation of the HVCM components over time. The correctness of operation can thus be verified.

HVCM FAULT DETECTION, DEGRADATION AND TUNING

The goals for applying ML to HVCM are twofold: (1) to reduce unscheduled downtime of the accelerator by predicting the degradation of the HVCM components, and (2) to reduce the tuning time and improve performance of the manual tuning of the modulators through automation and learning.

The modulators are currently tuned before turning over to operations to compensate for voltage droop as the capacitor bank discharges during the pulse, and to minimize ripple on the flat top. Since the high voltage circuit is a resonant circuit consisting of the transformer and its leakage inductance, and a resonant capacitor across the transformer secondary, the gain of the circuit is frequency dependent. We utilize this dependency to modulate the switching frequency of the IGBTs so that the gain increases as the capacitor droops. The frequency is linearly swept from lower to higher frequency during each pulse, and the settings for the sweep are tuned for each modulator independently by watching the voltage droop, IGBT peak and commutating currents and transformer core flux to ensure they all stay within acceptable ranges during the tuning process. Overshoot and ringing on the flat top are then minimized in a similar fashion by adjusting the firing sequences of the IGBT phases and the width of the initial pulses during the rising edge of the output pulse. The tuning is saved and not set again unless maintenance is performed on that modulator. This process should be fairly simple to automate, but would need supervision until some confidence is obtained since mistuning could result in catastrophic failures if an IGBT current is exceeded for example. Automation of the process would allow for near real time tuning to adjust for things like variations in the AC line voltage, or possibly retuning as components degrade with age to keep the modulator performing at an optimal tuning until a scheduled maintenance period.

HVCM availability has increased markedly over the years as improvements have been made and additional machine safety systems added such as desaturation detection for the IGBT to turn the switches OFF during a fault such as a core saturating. This is done in hardware, tripping the modulator OFF. These faults are detected and the IGBTs turned off in the microsecond time scale, and normally the modulator can return to operations after a reset. While we are investigating the use of ML on these types of faults, it is not clear what if anything could be done to prevent the modulator from tripping.

Several fault conditions which have resulted in significant downtime over the past few years are failures in the rectifier stacks caused by counterfeited diodes, long term degradation of the film-foil capacitors used in the resonant circuits, and polymerization of the FR3 insulating oil resulting in thermal runaway. Currently, we do not have the instrumentation for real time dissolved gas analysis needed to decide if the oil needs to be changed, however, the other faults listed above can be detected as shown in the SPICE analysis of Fig. 4. After training the autoencoder ML model on simulated data, we plan to test on one of the HVCM test stands at the SNS using different values of resonant capacitors.

Our first objective here is to predict upcoming faults so as to enable graceful shutdown and avoid catastrophic failure. We use anomaly detection techniques for fault prediction. A typical anomaly detection algorithm includes two steps. The first stage consists of data transformation to enhance the characteristic features for the nominal data. Example approaches include transforming the recorded waveform from the time domain to frequency domain or using ML models such as convolutional neural networks (CNNs) to capture hidden spatial/temporal structure from the source data. The second stage aims at anomaly detection. This step includes a definition of an anomaly measure and a determination of a threshold to decide under which conditions the

ISBN: 978-3-95450-214-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2021-THPAB252

waveform observed at a specific time represents an anomalous pattern. Recent deep learning model structures used for anomaly detection include variational autoencoders, generative adversarial networks, and normalizing flow. These deep learning models map different types of source data for anomaly detection application. They are data hungry and potentially computationally expensive and present some challenges for direct use for real-time fault detection/prediction in our HVCM application that has relatively small training data sets representative for each type of fault.

ML FOR FAULT PREDICTION

We propose a data-efficient anomaly detection technique to predict HVCM faults. Our approach will be analyzing the data on a moving window. We first transform all the waveforms characteristic collected in normal operation, in a given time window, from the time domain to the frequency domain using discrete cosine transform and calculate the mean, μ, and variance $\mathbf{q}$. We then utilize these values for predicting anomalies in incoming real-time measured data. Specifically, for a given waveform in the same time window, we first transform the data to the frequency domain obtaining frequency data $\mathbf{f}$ and we then calculate $\mathbf{p} = (\mathbf{f} - \mu)^2$ for each step in the moving window. Finally, we define the anomaly measure a as the $\|\mathbf{p} - \mathbf{q}\|_2 / \|\mathbf{q}\|_2$ where $\|\cdot\|_2$ represents the L_2 norm. Last, we determine a threshold λ for anomaly detection such that when $a > \lambda$ then an anomaly is detected. The determination of the threshold λ affects the performance of the anomaly detection method. Here we calculate the anomaly measure a for all waveforms in the nominal data set and choose the mean value of a as the λ by considering the normal data variation so as to balance the false negative and false positive rate.

We applied our anomaly detection and prediction method to B flux of one of the modulators (modulator SCL09 in Fig. 3). The available data consists of 83 B-flux waveforms with a sample rate of 400 *ns*, where 67 waveforms are labeled as *nominal* and 16 are labeled as *faulty*. Figure 5 plots the 83 waveforms in the 10000 time steps during the time inter-val 0.016-0.02 seconds. The waveforms labeled as nominal are presented in the top subgraph and their similarity is eas-ily observed. The 16 waveforms representative for faulty operation are presented in the bottom subgraph. We use the proposed anomaly detection method to provide advance notice of an upcoming fault. Figure 6 presents the prediction results of one of the 16 fault waveform. The blue curve is the recorded waveform; the red vertical line indicates the time when the fault occurs, and the green dashed vertical line identified the time when the fault is detected using our anomaly detection method.

Our method can predict the fault 1860 μs in advance that provides a sufficient time to take action for graceful shutdown and minimize the impact of upcoming failures. We successfully predict 11 out of 16 faults where the false negative rate is about 30% and we are improving our algorithm to reduce this rate. This preliminary result indicates that ML can provide advance notice of upcoming faults.

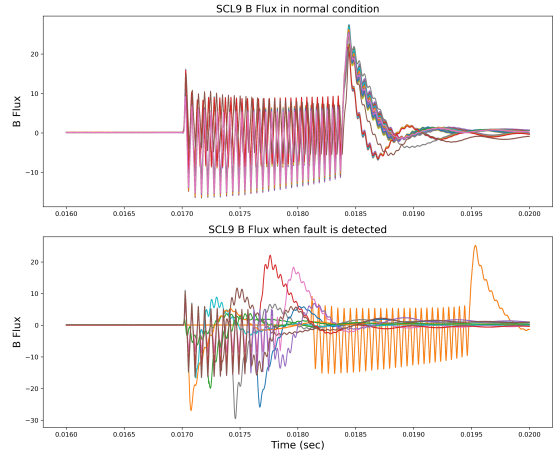

Figure 5: B flux data. Top: 67 nominal data. Bottom: 16 faulty data.

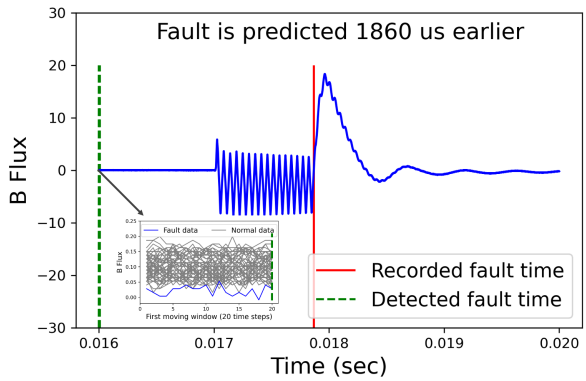

Figure 6: Anomaly prediction results. We use a window size of 20. Based on the first moving window of data (shown in the lower left box), our method detects the fault.

CONCLUSION

In this paper we discussed some opportunities for improving reliability of SNS operation and improvements in beam availability by predictive maintenance of HVCMs and reduced tuning and commissioning time of HVCM controls. We outlined ML approaches that will leverage simulation data and existing HVCM operational data for component health state prediction. We also introduced a fault prediction methodology that does not require large training data sets, and demonstrate its performance with data from one HVCM. Part of future work we intend to address the high miss detection rates by means of more sophisticated data transformation methods and by data augmentation including the use of simulated data. HVCM component health state estimation and HVCM autonomous re-tuning are next steps.

ISBN: 978-3-95450-214-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2021-THPAB252

REFERENCES

[1] W. A. Reass *et al.*, "Operational results of the spallation neutron source (sns) polyphase converter-modulator for the 140 kv klystron rf system," LANL, Los Alamos, NM, USA, Rep. LA-UR-01-3155, Jan. 2001.

[2] V. V. Peplov, D. E. Anderson, R. I. Cutler, M. Wezensky, J. D. Hicks, and R. B. Saethre, "SNS Linac Modulator Operational History and Performance", in *Proc. 24th Particle Accelerator Conf. (PAC'11)*, New York, NY, USA, Mar.-Apr. 2011, paper TUP275, pp. 1340–1342.

[3] M. Howell, B. DeGraff, J. Galambos, and S.-H. Kim, "SNS proton power upgrade," *IOP Conference Series: Materials Science and Engineering*, vol. 278, p. 012185, Dec. 2017. doi:10.1088/1757-899x/278/1/012185