ISBN: 978-3-95450-208-0 doi:10.18429/JACoW-IPAC2019-TUXXPLM2

SRF CAVITY FAULT CLASSIFICATION USING MACHINE LEARNING AT CEBAF*

A. Solopova†, A. Carpenter, T. Powers, Y. Roblin, C. Tennant, Jefferson Laboratory, Newport News, VA, USA
L. Vidyaratne, K. Iftekharuddin, Vision Lab, Department of Electrical and Computer Engineering, Old Dominion University, Norfolk, VA, USA

Abstract

The Continuous Electron Beam Accelerator Facility (CEBAF) at Jefferson Lab is the first large high power CW recirculating electron accelerator to make use of SRF accelerating structures. The structures are configured in two antiparallel linacs connected by arcs. Each linac consists of twenty C20/C50 cryomodules each containing eight 5-cell cavities and five C100 upgrade cryomodules each containing eight 7-cell cavities. Accurately classifying the source of cavity faults is critical for improving accelerator performance. A cavity fault triggers a waveform acquisition process where 17 waveform records sampled at 5 kHz are recorded for each of the 8 cavities in the affected cryomodule. The waveform record length is sufficiently long for transient microphonic effects to be observable. This data combined with archived signals sampled at 10 Hz are used to classify faults. Significant time is required for a subject matter expert to analyze and identify the intra-cavity signatures of imminent faults. This paper describes a path forward that utilizes machine learning for automatic fault classification. Post-training identification of the physical origins of faults are discussed, as are potential machine-trained model-free implementations of trip avoidance procedures. These methods should provide new insights into cavity fault mechanisms and facilitate intelligent optimization of cryomodule performance.

DEFINITION OF THE PROBLEM

The 12 GeV Upgrade for CEBAF was completed in September 2017. The project doubled the beam energy of the existing accelerator. To meet this energy goal, eleven new 100 MV cryomodules (called C100s) and RF systems were installed in 2013 (see Fig. 1) [1]. Currently the largest contributor to CEBAF downtime are beam trips caused by SRF cavities. During the last year there were an average of 6 RF trips an hour, accounting to roughly 15% of lost beam time per hour every day. To reduce the trip rate accelerating gradient of the cavity needs to be lowered, which means energy reach of CEBAF suffers.

The cavities in a C100 cryomodule have strong cavity to cavity mechanical coupling. When one cavity trips off, the Lorentz force detuning causes vibrations in the cavity string that are sufficient to trip other cavities. In order to avoid trips, the entire string is switched to self-excited loop mode (frequency tracking) when one of the cavities trips and others become unstable. This is also the default response for various other off normal conditions, which makes it difficult to determine which cavity initiated the cascade of faults [2].

When a cavities trips off, it disrupts delivery of the beam to the experimental halls. Correctly classifying which of several known fault mechanisms caused the cavity to trip provides valuable information to control room operators on how to treat the offending cavity and ultimately helps to maintain greater beam availability to users [3].

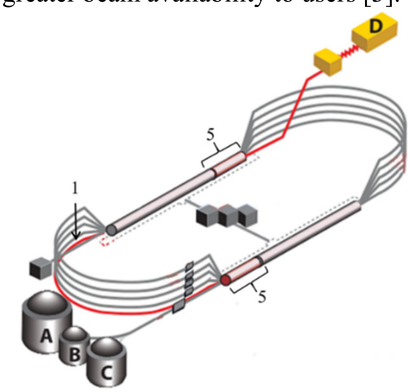

Figure 1: Schematic of the CEBAF accelerator showing the locations of the 11 C100 cryomodules from which cavity fault data is recorded.

USING THE FAULT IDENTIFICATION AND MACHINE OPERATION

Some examples that illustrate how prompt identification of fault types can be useful in machine operation:

- Fast Quenches: Identification of prompt "quenches" of cavities where the stored energy in the cavity is dissipated in times that are much shorter than is possible due to thermal quenches. These events were identified in the CEBAF operation as a gas discharge inside the cavity where stored energy is transferred to electrons produced by the discharge in times on the order of 10 μs. When these types of events occur in either the first or last cavity in the cryomodule there is pressure outburst observed in the beam line ion pump. In some cryomodules these events started occurring multiple times per day after weeks of no events and at gradients well below previously determined quench gradients. This can indicate gas loading in the beamline or the warm-to-cold transition in the RF waveguides. In addition to the temporary mitigation of reducing the gradient, identifying this type of fault can indicate a vacuum problem or the need to thermally cycle the cryomodule.

* Authored by Jefferson Science Associates, LLC under U.S. DOE Contract No. DE-AC05-06OR23177
† shabalin@jlab.org

ISBN: 978-3-95450-208-0 doi:10.18429/JACoW-IPAC2019-TUXXPLM2

- Thermal Quenches: The end groups of C100 cavities are conduction cooled to the 2K bath which surrounds the cavity cells. Slowly increasing temperature of the end group can lead to quenches in the cavities. This type of event is characterized by a single cavity having a major shift in frequency of about 45 Hz as measured closed loop phase difference between the incident and transmitted power signals. This phase shift only occurs in one out of the 8 cavities and has a time constant on the order of 100 ms. Faults such as this can be mitigated by reducing the cavity gradient slightly.
- Microphonics: these faults occur when there is a mechanical vibration of the cavities within the cryomodule. Such faults are characterized by the relative phase between the incident and transmitted power on all of the cavities oscillating in a coherent manner. Eventually one of the cavities faults due to insufficient RF power available to maintain a regulated gradient. Understanding that a fault is driven by microphonics rather than other mechanisms can help in determining if it is necessary to apply mitigations such as damping on external structures associated with the cryomodule or, for example, when multiple cryomodules have simultaneous microphonics trips that the source of the outside driving term needs to be mitigated.

DATA ACQUISITION

In order to better understand the nature and frequency of these faults, a waveform harvester has been implemented in the C100 cryomodules to record data for off-line analysis. For every cavity trip in a C100, the system automatically records 17 RF signals from each of the 8 cavities in the cryomodule. The digital LLRF system allows buffering of waveform data and saves ~1.6 seconds of data from the cavity trip event, including several hundred milliseconds of data prior to the event. An example of the online GUI is shown in Fig. 2. The recorded time-series data allows subject matter experts to analyze the faults in order to answer two questions:

1. Which of the 8 cavities within the cryomodule became unstable first?
2. What type of cavity fault caused the trip?

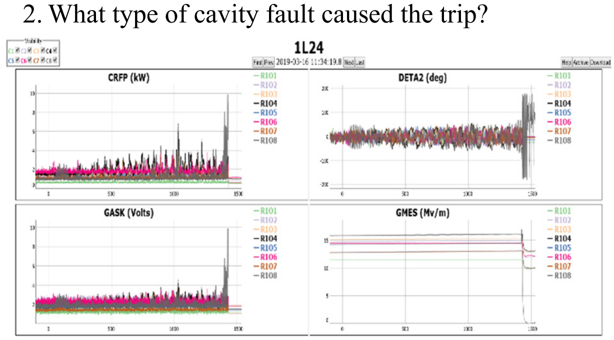

Figure 2: Online GUI for fault classification.

In our data set there are 5 types of faults: single cavity turn off, quench, prompt quench, microphonics and multi-cavity turn off. We ignore the last one because it describes an event when all cavities in the string turned off, which is usually associated with an external trigger such as vacuum valves closing or an intentional RF off command. Since all cavities trip simultaneously, this type of fault is of little use for answering one of our primary questions. Hundreds of labelled examples of RF faults from three separate accelerator run periods currently exist. These were analyzed and manipulated into a series of summary files containing a cryomodule and cavity ID, fault type, time stamp and the number of the cavity which went unstable first.

SUPERVISED LEARNING

The first approach to the problem was to use shallow machine learning algorithms which require a given set of features. To extract time-series characteristics, also called features, a Python package called *tsfresh* was used. Using its comprehensive mode, it computes nearly 800 statistical parameters per signal, which is computationally expensive. In an effort to reduce the size of the set a “select features” function was utilized to keep what it believes to be the most significant ones. This reduced the number of features to 7758 from 13498. Several standard pre-processing steps were applied, such as replacing non-numbers with appropriate values and scaling to make sure functions are well-behaved and the learning step of model training converges faster.

Model Selection

A variety of machine learning models were trained to identify which one performs best. Data was split into a training (70%) and test (30%) set. A variety of classification models were trained, including k-Nearest Neighbours, Decision Tree, Support Vector and Gaussian Naive Bayes as well as ensemble models such as the Bagging Classifier, Random Forest, Extra Trees and Gradient Boosting. Fine-tuning the trained Decision Tree returns an accuracy score of 96.63% when applied to the withheld test data set [4]. Additionally, Autoregressive (AR) time-series feature extraction method was tested on select 5 signals. With AR features, several machine learning models were trained and evaluated using 10-fold cross-validation. The Random Forest classifier performed best achieving 91.5% accuracy. Same technique applied to all signals returns 88% accuracy for the Random Forest.

Feature Importance

One of the critical steps is the selection of important features. The Decision Tree classifier can rank the importance of each feature and the results are shown graphically in Fig. 3. Computing only the top 10 features greatly reduces

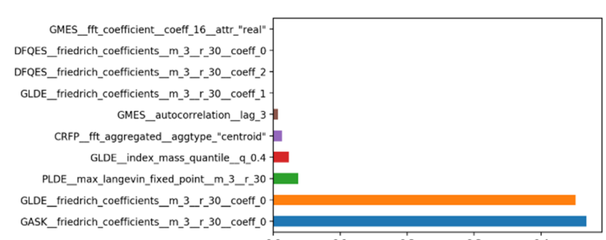

Figure 3: Feature importance from the Decision Tree Classifier.

ISBN: 978-3-95450-208-0 doi:10.18429/JACoW-IPAC2019-TUXXPLM2

the required computation time while maintaining good accuracy. An accuracy score of 94.38% using the Random Forest classifier was achieved using only the top 10 features [5].

First Unstable Cavity Identification

To determine which cavity tripped first, we follow the same process as above, but for each event we read in the 17 signals for all 8 cavities of the cryomodule (rather than just the cavity which tripped). To reduce the computational load, we used *tsfresh* to calculate features on only 3 of the signals (selected by subject matter experts) for each cavity. Training the same models, we obtain classification accuracies shown in Table 1 in which a trained model is applied to unseen test data [5]. No hyperparameter tuning was performed.

Table 1: Accuracy of Several Machine Learning Models Applied to Identifying Which Cavity Tripped First

Method	Accuracy Score (%)
Random Forest	95.72
Gradient Boosting Classifier	95.19
Extra Trees	94.12
Bagging Classifier	94.12
Support Vector Classification	90.91
k-Nearest Neighbor	88.77
Decision Tree	87.70
Gaussian Naïve Bayes	86.10

DEEP LEARNING

Shallow machine learning seemed to be an obvious first step for our problem, given the limited data set. However, other projects have used deep learning with good results with similar sized data sets [6]. As a test case, we implemented a deep Recurrent Neural Network (RNN), specifically a Long Short Term Memory (LSTM) architecture. This type of neural network is often used in time-series analysis where the temporal features of the data is important. To reduce training time we used only 5 of the 17 signals – the same ones used by experts for manual identification. A 10-fold cross-validation accuracy score of 86% was achieved, which we expect to increase with the number of samples and hyperparameter tuning. Deep learning avoids the feature engineering step and brings us closer to our ultimate goal of using the raw data in a real-time control room application.

FUTURE WORK

This paper represents only preliminary results, much work remains to complete the project. A recent collaboration with Old Dominion University will provide additional resources and expertise in machine learning to help solve the problem of identifying the first unstable cavity and, potentially, trip predecessors. Applying deep learning for fault classification could help avoid a human error by eliminating the need for feature selection and provide more accurate data for cavity identification. Our next goal is to move from off-line analysis to a near real-time diagnostic system.

Trip Avoidance

When a cavity fault occurs due to insufficient RF power, for example a microphonics related fault, one possible avoidance mechanism is to momentarily (10-20 ms) increase the available power. For CEBAF continuous power supplies, this is not currently possible. However, the addition of an extra pulsed power supply may be an option. This could be implemented as an anomaly detection problem if we are able to collect "good" data and monitor cavity conditions in real time [7]. A similar approach can be taken to compensate for mechanical vibrations using piezo-tuners. This must be done with care, since C100 cavities are strongly coupled, and increasing bandwidth may result in piezo-tuners causing more vibrations and instabilities. The potential of piezo tuners to compensate for vibration was demonstrated on new LCLS-II cryomodules during acceptance testing at Jefferson Lab [8].

CONCLUSION

Over the last year waveform data has been collected from multiple CEBAF run periods. We started applying machine learning techniques in late November 2018 and have made significant progress over the last few months (see Table 2.)

- Trained and evaluated several ML models, achieving 96.6% accuracy after fine-tuning the Decision Tree model
- Using only the top 10 features from the Decision Tree model, we have obtained 94.38% accuracy with a Random Forest Classifier (no hyperparameter tuning).
- Performed autoregressive feature extraction on raw data, achieving 91.5% cross-validation accuracy with random forest classifier
- Implemented a LTSM deep neural network on raw data with a promising 86% accuracy score.
- Achieved 95.72% accuracy using a Random Forest model (no hyperparameter tuning) to classify which cavity tripped first.
- Understanding the origins and the nature of cavity faults is important for increasing accelerator reliability and availability as well as for improvement of the future cryomodules design. Timely delivery of the information to the control room is essential.

Table 2: Results of SRF Cavity Fault Classification Using Machine Learning

	Machine Learning			Deep Learning
Problem	Which Cavity?	Which Fault?		Which fault
Feature engineering	*tsfresh*	*tsfresh*	auto-regression	none
Model	Random Forest	Decision Tree	Random Forest	RNN-LSTM
Result	95.70%	96.60%	88%	86%

ISBN: 978-3-95450-208-0 doi:10.18429/JACoW-IPAC2019-TUXXPLM2

REFERENCES

[1] M. F. Spata, "12 GeV CEBAF Initial Operational Experience and Challenges", in *Proc. 9th Int. Particle Accelerator Conf. (IPAC'18)*, Vancouver, Canada, Apr.-May 2018, pp. 1771-1775. doi:10.18429/JACoW-IPAC2018-WEYGBD1

[2] T. Powers, C. Hovater, A. Solopova, "C100 Cavity Faults as Determined by Time Domain Waveforms", presented at TTC Workshop, Canada, 2019.

[3] Y. Roblin, C. Tennant, A. Solopova, "Machine Learning Applications for the CEBAF Accelerator", JLAB-TN-19-017.

[4] C. Tennant, T. Powers, Y. Roblin, A. "Recent Results of SRF Cavity Fault Classification Using Machine Learning at Jefferson Laboratory", presented at the 2nd ICFA Workshop for Machine Learning for Particle Accelerators, Switzerland (2018).

[5] C. Tennant, A. Carpenter, T. Powers, Y. Roblin, A. Solopova "Progress on Using Machine Learning for SRF Cavity Fault Classification", JLAB-TN-19-018.

[6] J. Edelen, private communication (2019).

[7] C. Hovater, R. Nelson, private communication (2019).

[8] C. Hovater *et al.*, "Commissioning and Operation of the CEBAF 100 MeV Cryomodules", in *Proc. 3rd Int. Particle Accelerator Conf. (IPAC'12)*, New Orleans, LA, USA, May 2012, paper WEPPC093, pp. 2432-2434.