ISBN: 978-3-95450-232-5 ISSN: 2673-7000 doi:10.18429/JACoW-NAPAC2022-WEPA40

THE L-CAPE PROJECT AT FNAL

M. Jain*, J. Strube[1], V. Amatya, G. Panapitiya
Pacific Northwest National Laboratory, Richland, WA, USA
B. Harrison, K. J. Hazelwood, W. Pellico, B. Schupbach, K. Seiya, J. St. John
Fermi National Accelerator Laboratory, Batavia, IL, USA
[1]also at University of Oregon, Eugene, OR, USA

Abstract

The controls system at FNAL records data asynchronously from several thousand Linac devices at their respective cadences, ranging from 15 Hz down to once per minute. In case of downtimes, current operations are mostly reactive, investigating the cause of an outage and labeling it after the fact. However, as one of the most upstream systems at the FNAL accelerator complex, the Linac's foreknowledge of an impending downtime as well as its duration could prompt downstream systems to go into standby, potentially leading to energy savings. The goals of the Linac Condition Anomaly Prediction of Emergence (L-CAPE) project that started in late 2020 are (1) to apply data-analytic methods to improve the information that is available to operators in the control room, and (2) to use machine learning to automate the labeling of outage types as they occur and discover patterns in the data that could lead to the prediction of outages. We present an overview of the challenges in dealing with time-series data from 2000+ devices, our approach to developing an ML-based automated outage labeling system, and the status of augmenting operations by identifying the most likely devices predicting an outage.

INTRODUCTION

Machine Learning (ML) and Artificial Intelligence (AI) have become ubiquitous in recent years, with applications ranging from computer vision to speech recognition to increasing the autonomy of controlled systems. This development has in part been furnished by the proliferation of data sources: Cheap, easy to deploy sensors permit the instrumentation and monitoring of ever more systems. The controls system for the accelerator complex at the Fermi National Accelerator Laboratory (FNAL) monitors and issues commands to 4000+ control system parameters in the linear accelerator (Linac) at frequencies ranging from 15 Hz to once every few minutes. Upon the start of a beam interruption, this data is used by accelerator operators to investigate the source of the unplanned beam outage from the FNAL Main Control Room. The electric power consumed during these outages adds up to a considerable amount of energy used (see Table 1) and could be reduced significantly if the duration and type of outage were identified quickly and accurately. We present a pipeline to enable the use of ML/AI that augments the data flow to the control room with analytics of outages, reducing the time to label them meaningfully and minimizing the number of incorrect or inconsistent labels.

* milan.jain@pnnl.gov

Table 1: Energy Use During Beam Outages

Event Length	Event Count	Duration (hrs)	Energy (MWh)
< 5 s	1626	1.04	0.182
< 10 s	321	0.59	0.104
< 60 s	205	2.01	0.351
< 120 s	201	4.70	0.822
< 10 m	169	11.69	2.05
≥ 10 m	47	111.35	19.49

DESCRIPTION OF THE DATA SOURCES

The accelerator control system's Data Logger nodes record data streams into circular buffers. To store this data for a longer period than the lifetime of the circular buffers, this project requested a data acquisition pipeline to write the data to long-term storage using an industry standard format, HDF5 [1]. This pipeline was created by the Controls Department developers using modern tools to solve a common problem and is being used on other ML projects. This new pipeline allows projects to choose a data source without modifying requests allowing L-CAPE to switch from using historical data to live data seamlessly.

In the current configuration, L-CAPE makes 5567 requests over 4292 control system parameters and stores each request in an HDF5 group. The HDF5 output is collected by the hour with an average file size of 644 MB per period.

DATA PREPROCESSING PROCEDURES

Hourly raw data in HDF5 format is sampled from devices at their respective cadences and timestamped by independent front-end nodes' clocks. A reference clock is required to time align the data for analyzing and modeling the devices together. By using a reference clock capturing the highest frequency devices (which is 15 Hz in this case), loss of information can be avoided for all devices. Therefore, for time alignment, we use timestamps at 15 Hz starting at the first timestamp of each hourly data file. Because there is a new file every hour, the reference clock reset at the end of every hour, thus reducing shift in data over time because of 15 Hz sampling rate.

Using a reference clock and combining the data removes redundant timestamps, which, together with applying the lossless snappy compression algorithm [2], reduces the disk space taken by the data (~20x), accelerates data read/write operations, and allows practical analysis and visualization of multiple devices across many days simultaneously. Furthermore, for more speedup in read/write, this combined data is

ISBN: 978-3-95450-232-5 ISSN: 2673-7000 doi:10.18429/JACoW-NAPAC2022-WEPA40

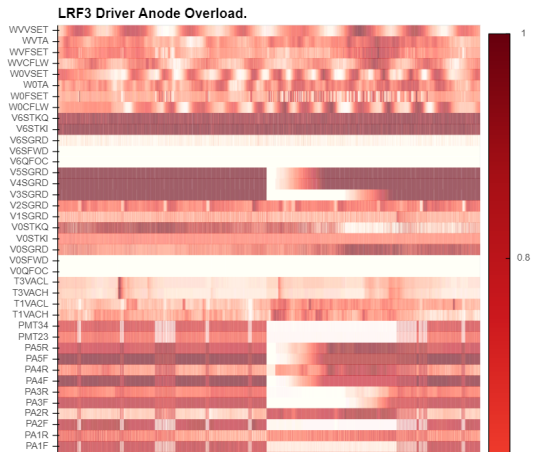

Figure 1: Heatmap visualization of a fault.

Table 2: Evaluation of Operator-Defined Labels

	True Positive	False Negative	False Positive
Bit 27	58	54	161
Bit 28	64	48	111
27 | 28	84	28	448
27 & 28	25	87	7

stored in a binary Parquet [3] file. For analysis and visualization, the data is loaded into a Pandas DataFrame [4, 5]. Figure 1 shows the heat map visualization of the data (normalized range, 0–1) from a subset of devices for a single fault instance – "LRF3 Driver Anode Overload". Different devices show different patterns during this fault, with varying start time and duration. On comparing this heat map with other faults and other instances of the same fault, we found that the pattern could vary not only from one fault type to another but also from one instance of a fault to another.

OUTAGE LABELING

First stage of our workflow is to train a model that captures "normal" operation. Thus, we created a list of outages with labels and times reported by the operators. A data-driven method to identify downtimes could use the beam permits (bits 27 and 28 for the downstream and upstream permits, respectively), but given the challenge of identifying a combination that agrees with the human-created labels (see Table 2 for an agreement between different combinations of the permit bits and the operator-defined labels), we proceed with the latter as the "ground truth" for this study.

MODELING

Since training of neural networks for all devices is compute expensive, we explored two statistical modeling techniques (thresholding and filtering) in addition to an ML-based method. Thresholding assumes that device value lower than the 5th percentile or higher than the 95th percentile is anomalous. Thus, for thresholding, the reconstruction error at any timestamp is defined in Eq. 1, where R_e^t and D^t indicate the device value at time t, and D_{low} and D_{high} depict 5th and 95th percentile from the device's data distribution:

$$R_e^t = \begin{cases} D_{\text{low}} - D^t, & D^t < D_{\text{low}} \\ 0, & D_{\text{low}} \leq D^t \leq D_{\text{high}} \\ D^t - D_{\text{high}}, & D_{\text{high}} < D^t \end{cases}. \quad (1)$$

In another statistical approach, the data is passed through a low-pass filter (to filter the high-frequency noise) followed by a high-pass filter that filters brief anomalies/spikes. The filtered signal is then compared with the actual signal to compute the reconstruction error.

Third, we implemented TADGAN - an unsupervised anomaly detection approach built on Generative Adversarial Networks (GANs) by Geiger *et al.* [6]. TADGAN can be trained on multivariate time series data, however, given the dynamic logging of devices and difference in sampling frequency of the devices, training a single model for all devices was infeasible. Therefore, one model per device was trained and the reconstruction error for each device is the difference between the reconstructed signal and the actual signal.

For all three techniques, the final anomaly score is the sum of device-level reconstruction error. Since data from devices is sampled at different frequencies, before calculating the anomaly score, device-level errors are downsampled to 1 s with aggregator function `max`. Figure 2 compares anomaly score from TADGAN with operators' label (yellow region) and bit permits (dotted lines) on a specific day.

EVALUATION

Table 3 summarizes the detection accuracy based for one month of data (Mar-2020) for all three techniques. Based on the feedback from the domain experts, devices related to status bits and control parameters were filtered for the modeling, thus limiting the device count to 2081. Since training TADGAN is compute expensive, only 53 devices with complex data patterns were trained using the TADGAN. If compared on f1-score, the analysis indicates that statistical techniques are performing better (high recall) than the TADGAN. However, if the distribution of device data is multi-modal (for instance, devices running at different settings), simpler statistical techniques like thresholding tends to perform poor. Akin to that, while filtering performs well on longer time-series, it tends to perform poorly on smaller time windows, and thus for real-time inference. Therefore, an ensemble of both statistical and deep-learning models are required for a reasonably accurate fault detection framework.

When compared with downstream (bit-27) and upstream bit permit (bit-28), we noticed that the false negative increases significantly. Figure 3 explains the reason for this for the filtering approach. The fault duration plot (on the right) shows bit permits frequently go down only briefly, and such instances are neither reported in the ground truth, nor can be detected by any of the techniques discussed.

Once a fault is detected, the next step is to explain and label it. The operator's labeling depends on subjective opinion and experience, and therefore could be rather inconsistent.

ISBN: 978-3-95450-232-5 ISSN: 2673-7000 doi:10.18429/JACoW-NAPAC2022-WEPA40

Table 3: Performance Comparison. Data for March, 2021

	#Devices	True Positive	False Negative	False Positive	Precision	Recall	F1-Score
Thresholding	2081	99	13	210	0.32	0.88	0.47
Filtering	2081	105	7	279	0.27	0.94	0.42
TADGAN	53	51	61	443	0.1	0.45	0.17

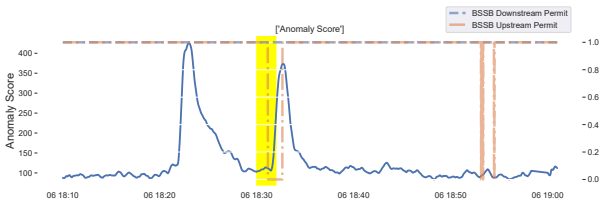

Figure 2: TADGAN anomaly score for different pathologies.

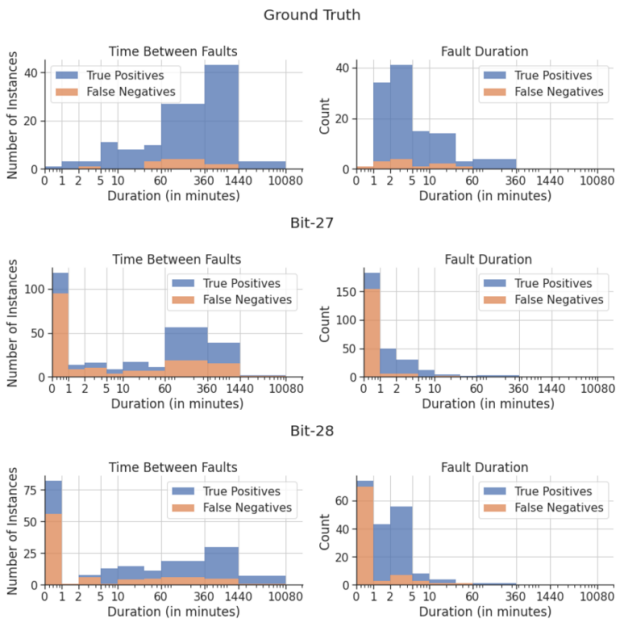

Figure 3: Evaluating the false negatives.

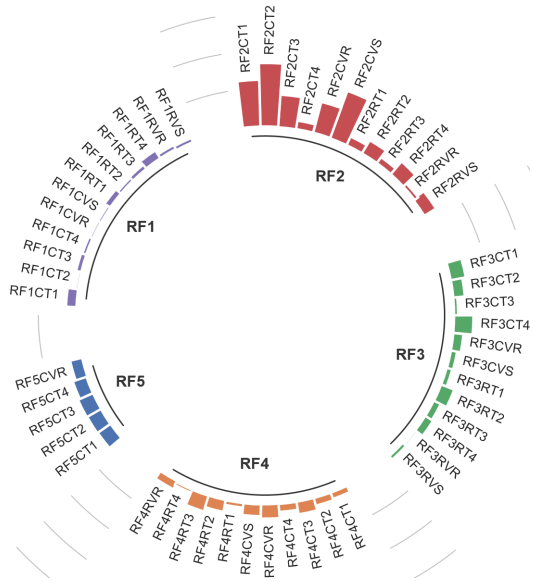

(a) LRF2 Driver Anode OL

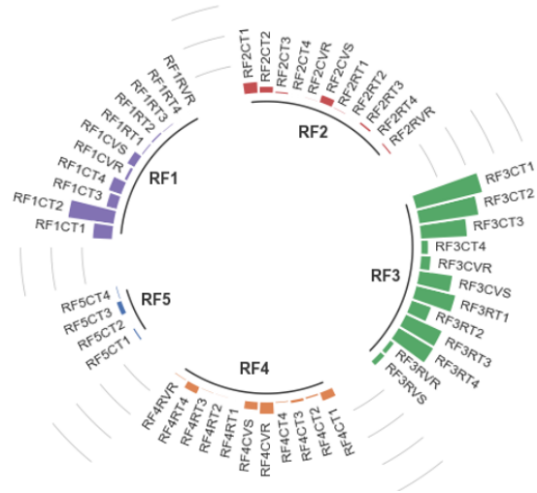

(b) LRF3 Driver Anode OL

Figure 4: Different faults can be distinguished by their signatures in different subsystems (RF1-5).

In our preliminary analysis, we noticed that faults correspond to a specific pattern of device-level reconstruction errors. Figure 4 compares two such faults with the help of a circular bar chart, where bar length indicates mean of max reconstruction errors during all the instances of two faults. It is evident that devices with high reconstruction error are correlated with the ground truth label.

CONCLUSION AND OUTLOOK

We have presented a pipeline to process data from the FNAL accelerator control system for automated fault detection and labeling. Our tailored approach of applying simple thresholding and filtering methods to model normal behavior by default and using machine learning techniques only where necessary provides a good balance of computational throughput and algorithmic performance. The high dimensionality of the data provides adequate separation of different types of outages. Our pipeline of data cleaning, alignment, and basic analysis is currently being implemented for near-real-time data analytics that augment the information available in the control room. Future studies will improve the performance of the labeling and prediction, and the data throughput to improve the decision-making ability in the control room with the goal of achieving energy savings by improving the operational efficiency and predictive capability.

ACKNOWLEDGMENTS

This work was supported in part by the U.S. DOE Office of Science, Office of High Energy Physics, under award 76651: "Machine learning for Accelerator Operations". Pacific Northwest National Laboratory is operated by Battelle Memorial Institute for the U.S. Department of Energy under Contract No. DE-AC05-76RL01830. It was also partly supported by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

ISBN: 978-3-95450-232-5 ISSN: 2673-7000 doi:10.18429/JACoW-NAPAC2022-WEPA40

REFERENCES

[1] The HDF Group, *Hierarchical data format, ver. 5*, 2022. https://www.hdfgroup.org/HDF5

[2] Google: *Snappy, a fast compressor/decompressor*, 2022. https://github.com/google/snappy

[3] D. Vohra, "Apache parquet," in *Practical Hadoop Ecosystem: A Definitive Guide to Hadoop-Related Frameworks and Tools*. 2016, pp. 325–335. doi:10.1007/978-1-4842-2199-0_8

[4] W. McKinney, "Data Structures for Statistical Computing in Python," in *Proc. 9th Python in Science*, Austin, TX, USA, Jun. 2010, pp. 56–61. doi:10.25080/Majora-92bf1922-00a

[5] J. Reback *et al.*, *Pandas-dev/pandas: Pandas v1.4.2*, 2022. doi:10.5281/zenodo.6408044

[6] A. Geiger, D. Liu, S. Alnegheimish, A. Cuesta-Infante, and K. Veeramachaneni, "Tadgan: Time series anomaly detection using generative adversarial networks," in *Proc. Big Data'20*, Virtual, Oct. 2020, p. 33. doi:10.1109/BigData50022.2020.9378139