Anomaly Detection for Equipment Failure at BNL

Ian Blackler, Operations Coordinator, Operations Group
Yuan Gao, Assistant Scientist, Controls Group

ACCELERATOR RELIABILITY WORKSHOP
ARW 2022

Introduction

The Collider-Accelerator Complex (C-AD) at Brookhaven National Lab spans over 2 square miles and comprises thousands of different elements, and tens of thousands of control values. Monitoring tolerances for these systems can be challenging. Usually systems self-report output values, performance data, and on/off status. These reports are delivered to the Operations staff through specialised software. However, there are some systems that provide none of this feedback. For example, one concern is the failure of facility systems like Air Conditioning units that have no connection to the C-AD complex. Catching failures like these is a high priority as they can lead to extensive downtime. Hence a machine-learning approach has been considered.

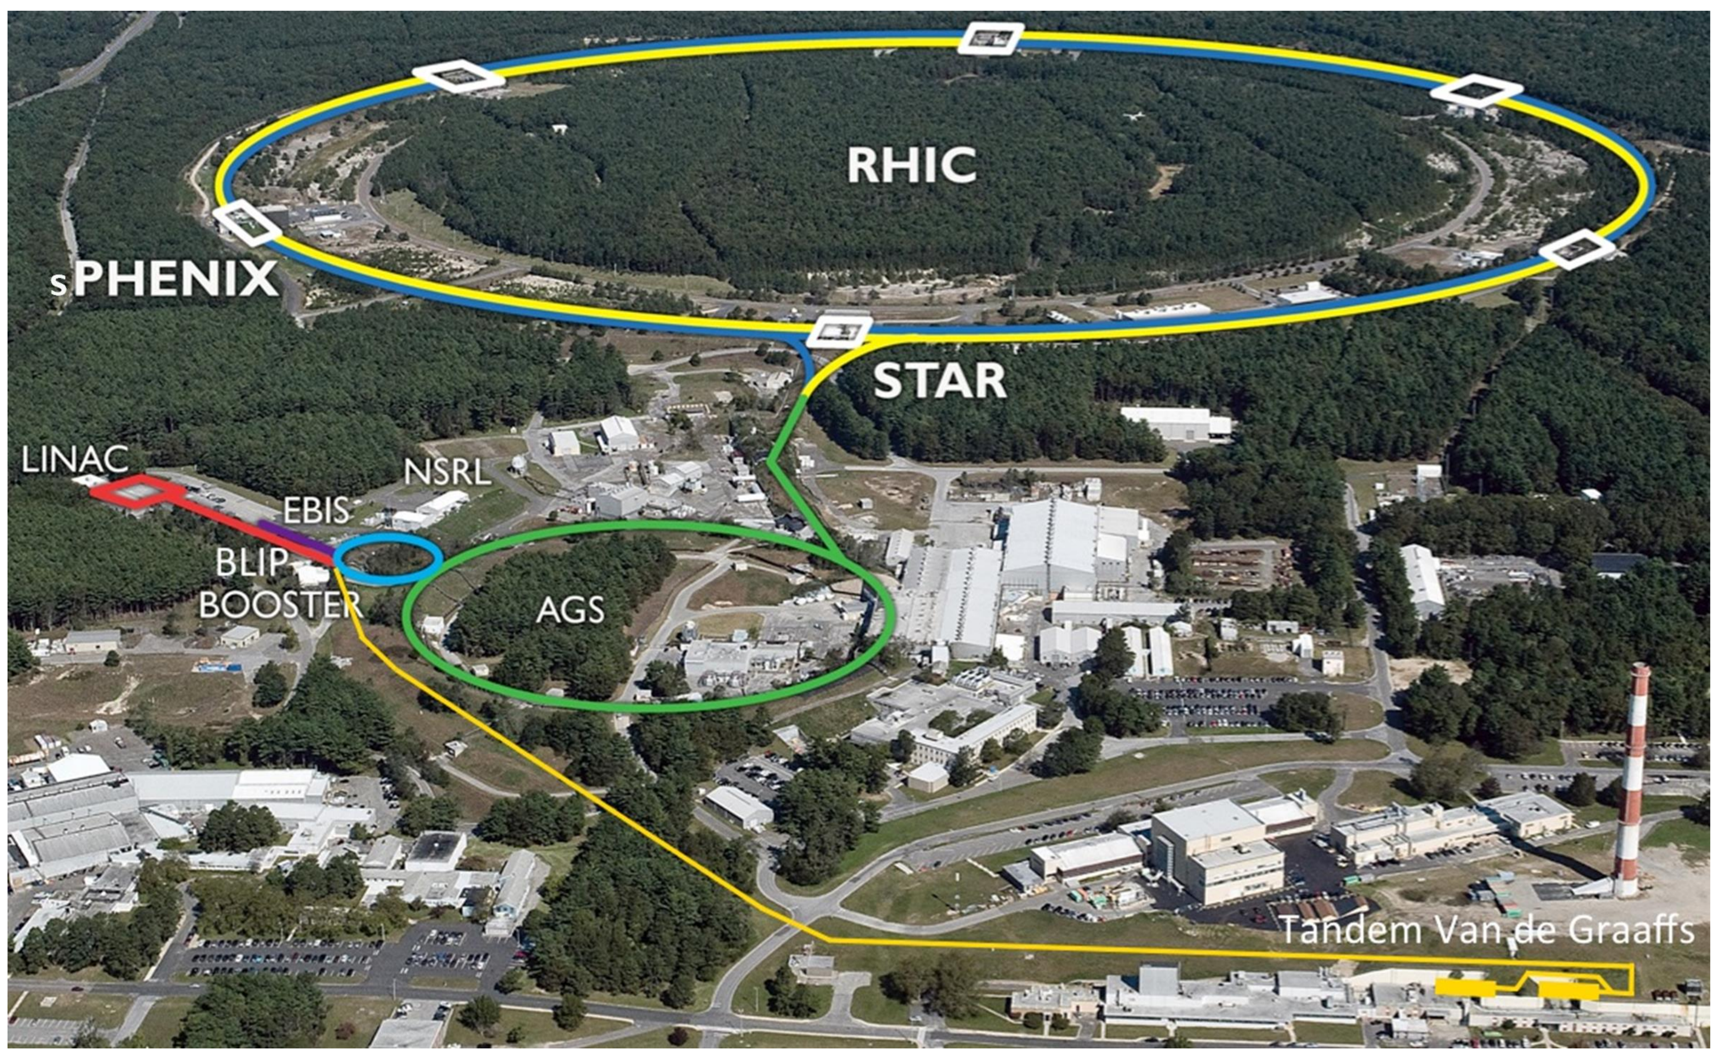

Secret Failures

In this example, an air-conditioning device failed by icing up. The room temperature is observed to drop rapidly, then rise steadily to alarm levels. By the time alarm limits had been reached, the system had been in failure for approximately 9 hours, and the technicians had been gone for 3 hours.

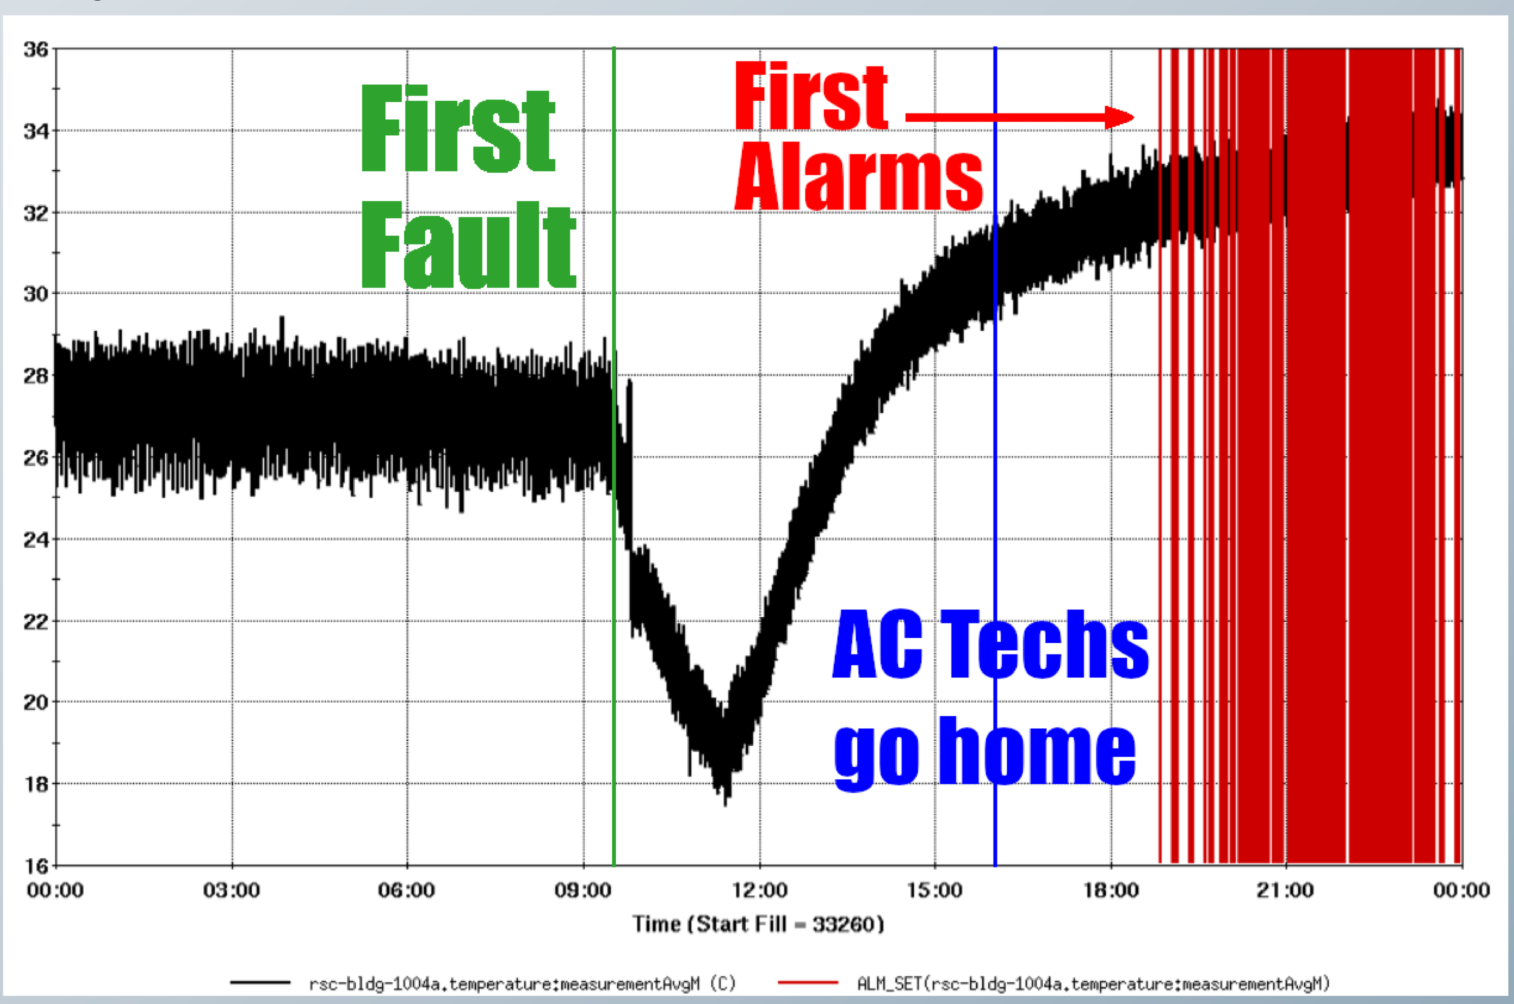

Training the Autoencoder

Detecting "secret" faults when they occur, or even predicting them before they occur is clearly a high value proposition. We started with sample data from a thermometer in the High Bay of Building 1012A of the RHIC Ring.

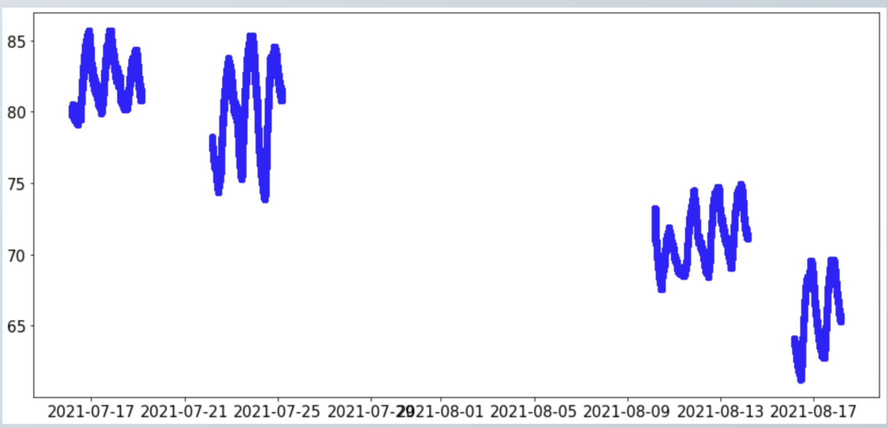

This is 1Hz data for several days in July-August 2021. This yielded 1,036,750 data points from four well-behaved sets.

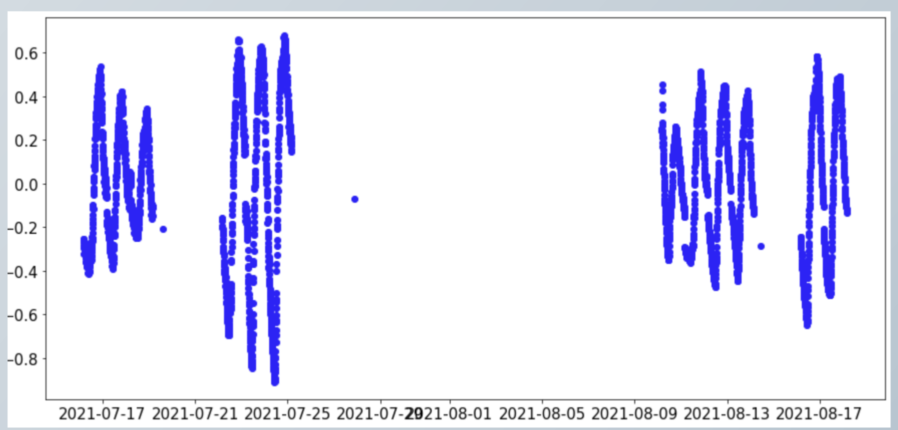

The data sets were then reduced to 1 average measurement every 300 seconds. This improved training time efficiency tremendously.

The data was then normalised and shifted to the centre using a 24-hour shifting window.

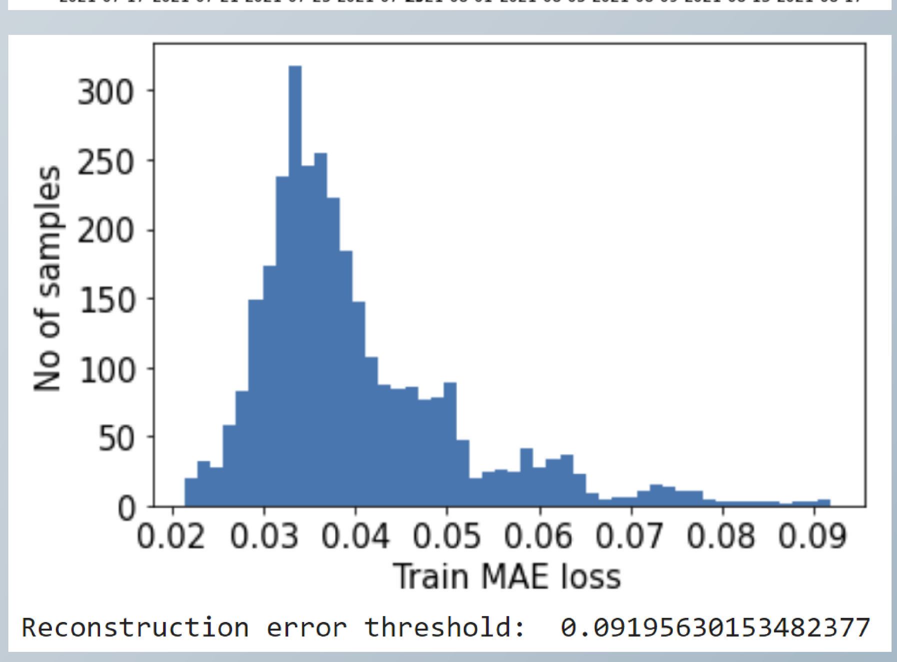

The distribution of MAE (Mean Absolute Error) values shows good performance with consistently low levels of mismatch.

Testing the Autoencoder

Some faults were observed in the A/C systems at Building 1012A. These lead to temperature anomalies that could be used to test the responsiveness of the trained autoencoder.

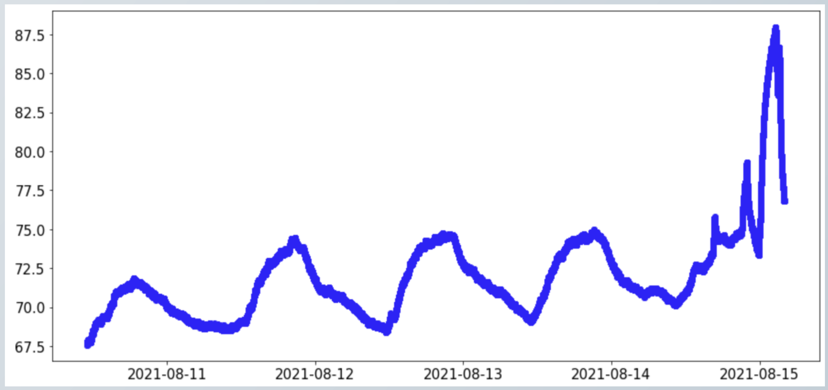

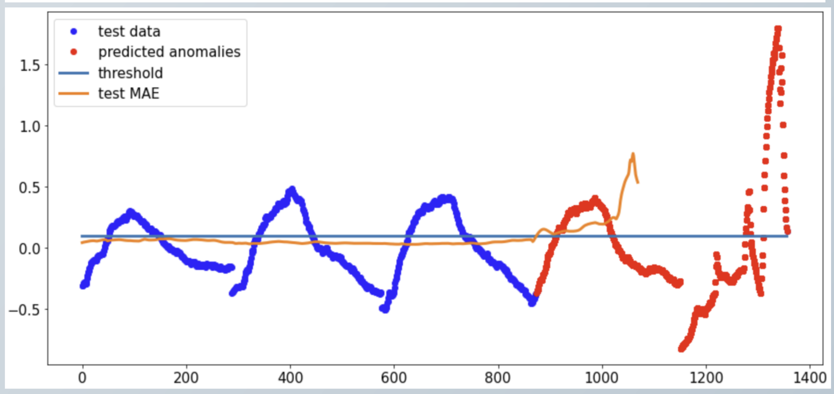

Here the test MAE is observed to respond remarkably well. Human observation would suggest that a fault occurred late on August 14th or early August 15th, but the MAE test signal begins to move far sooner.

Official alarms were not generated until 21:33, long after most experienced technical support had left for the day.

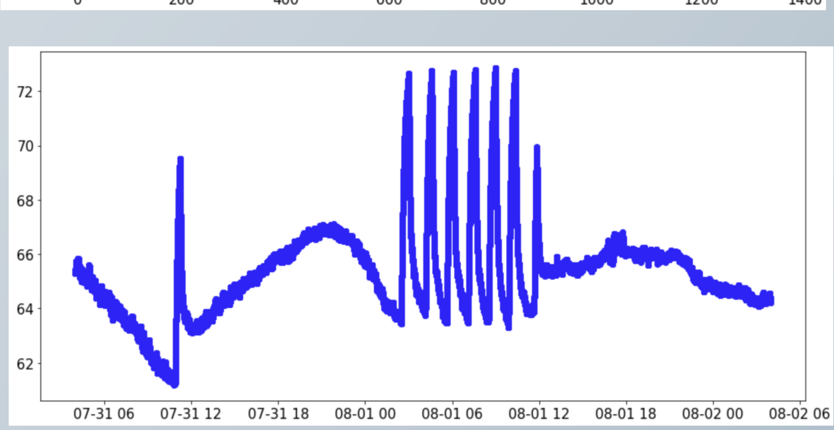

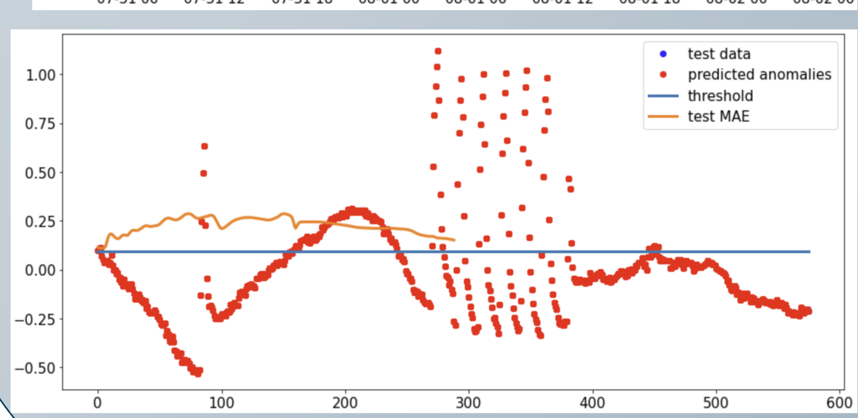

Here the test MAE showing problems long before any are observed in the alarm system.

The rapid cycling seen in the centre of the dataset is a result of the system being put into a non-standard mode of operation.

There is an open question to be addressed regarding how situations like these should be addressed.

Future Developments

A more sophisticated and complex approach could be taken to catching early anomalies, but C-AD has other concerns to consider. Firstly, we want this system to run live—this is not for post-mortem analysis—and we want it to run with as few false positives as possible. The little boy who cried wolf gets ignored by the Operators. Secondly, there are literally hundreds of temperature readbacks to train on. Each one behaves subtly—or conspicuously—different to the rest. Some track closely with outside temperatures, some have rapid cycling A/C units, some have massive heat loads that change according to the machine configuration. One size does not fit all. Finally, we want to examine generalising this approach to other systems, and to predicting failures before they occur.

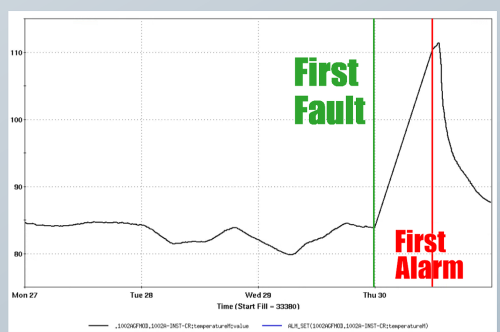

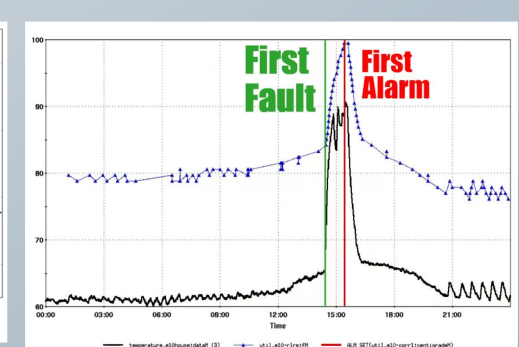

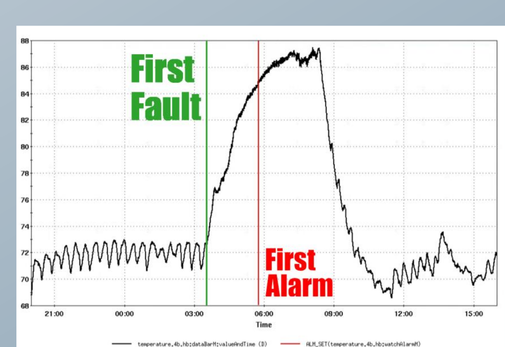

U.S. DEPARTMENT OF ENERGY

www.bnl.gov

Brookhaven National Laboratory