LEVERAGING LOW-COST SENSORS AND MACHINE LEARNING FOR PUMP ANOMALY DETECTION IN ACCELERATOR FACILITIES*

R. Sainju†, O. Mohsen, J. Ackley, J. Fuerst, R. Wright, Y. Sun, M. Borland
Advanced Photon Source, Argonne National Laboratory, Lemont, IL, USA

Abstract

The reliability of cooling systems is critical for removing substantial amounts (in megawatts) of waste heat from numerous high-power accelerator components (e.g., magnets, RF structures, power supplies) and beamline components. Reliance on manual inspection of hundreds of pumps is inefficient and increases the risk of costly component damage and unplanned downtime. This study introduces a real-time method for automated vibration-based pump anomaly detection that is designed to help transition facility operations from reactive to predictive maintenance. Our approach integrates (i) affordable vibration sensors and (ii) machine learning models for anomaly detection. We have deployed vibration sensors on Bunch Lengthening System (BLS) helium pump, and a set of deionized water-cooling pumps, where Linux edge nodes aggregate and preprocess the hourly collected data. To support operational decision-making, a web-based diagnostic platform provides real-time visualization of vibration trends against a rolling weekly baseline, together with Short-Time Fourier Transform (STFT) spectrograms and integrated Power Spectral Density (PSD) plots for frequency and time-domain analysis. Additionally, the web platform will display hourly inferences from the machine learning models such as autoencoder on the data stream, autonomously detecting spectral anomalies indicative of mechanical faults. The integration of scalable edge data collection, advanced visualization, and unsupervised machine learning will provide a vital safeguard for maintaining operational readiness in particle accelerators.

INTRODUCTION

Cooling water systems are an unglamorous but indispensable part of any modern accelerator. At APS, deionized water removes heat from magnets, radio-frequency structures, power supplies, and beamline components through a primary distribution system that operates at flow rates approaching 10,000 gallons per minute, together with dozens of smaller closed-loop secondary systems [1]. Pumps are the active elements that keep this network running, and an unexpected pump failure can cascade quickly: components interlock on flow and temperature, downtime is expensive, and damage to high-power loads is difficult to repair on a user-driven schedule.

Today, pump health at APS is assessed primarily through walk-down inspections, supplemented by reactive trips after the fact. Across hundreds of pumps this is labor-intensive, occurs periodically, and makes it difficult to catch the slow drifts that precede a hard failure. This bottleneck can be addressed through vibration-based condition monitoring [2], and by applying machine learning methods that learn the normal operations signature of a machine from unlabeled operating data [3, 4]. Recent work has shown that autoencoder-style reconstruction analysis transfers well to the accelerator setting: it has been applied to magnet fault prediction in the APS storage ring [5, 6] and to errant-beam and pressure anomalies at other facilities [7, 8]. Pumps, however, have received comparatively little attention in the accelerator literature. This work extends our earlier proof-of-concept on cooling water pumps [9] and adds the BLS helium pump and an web-based diagnostic platform.

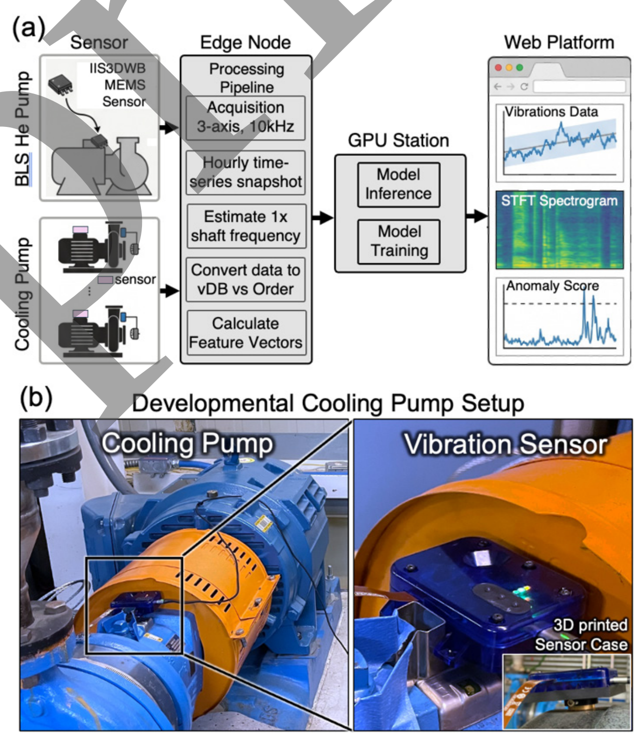

Figure 1: Developmental pump-monitoring system at APS. (a) End-to-end data flow: 3-axis MEMS vibration sensors mounted on accelerator pumps stream data to Linux edge nodes, which perform hourly aggregation, preprocessing, and feature vector calculations; the data is then sent to a GPU workstation for inference; outputs are served through a web-based diagnostic platform to engineers. (b) Representative sensor installation on the cooling pump, showing sensor mounting location, cable run to the local edge node, and 3D printed sensor case.

* Work supported by the U. S. Department of Energy, Office of Science, under Contract No. DE-AC02-06CH11357.
† rsainju@anl.gov

In this contribution, we summarize the design, deployment, and early experience of a developmental pump anomaly-detection system at APS. The system pairs an inexpensive 3-axis MEMS vibration sensor from STMicroelectronics [10] with on-site edge aggregation, hourly data acquisition and model inferences, and a web-based diagnostic platform that combines real-time data trends and anomaly detection alarms for human interpretation. We focus here on the system architecture and on the engineering lessons that have gathered while building this pump-monitoring system in a working accelerator environment; a more detailed treatment of the model architecture, training procedure, and quantitative detection performance is the subject of a separate publication in preparation.

DEVELOPMENTAL SETUP

Vibration Sensors and Edge Nodes

The vibration front-end utilizes the STMicroelectronics IIS3DWB, a three-axis MEMS accelerometer tailored for industrial vibration detection. It offers a flat frequency response up to about 6 kHz, features a steep anti-aliasing roll-off, low noise levels, and a digital SPI interface [10]. This bandwidth effectively captures the frequencies of the shaft rotation frequency and its lower-order harmonics, as well as the impeller and bearing characteristic frequencies that are diagnostic for centrifugal pumps and compressors. The sensors are mechanically coupled to the pump housing near the bearing block.

Figure 1 summarizes the end-to-end data flow. As shown in Fig. 1a, each sensor cluster is read out by a Linux edge node. Three-axis acceleration is acquired at for a fixed-length (30-seconds) segments and then passed through a preprocessing stage that produces both raw time-series snapshots and a set of vibration descriptors used downstream. As shown in Fig. 1b inset, we designed a 3D printed sensor case to mount a sensor to the pump. Time-series snapshots are retained on an hourly cadence so that models can be retrained from real operating data when needed, while the reduced descriptors are stored continuously to keep storage and network bandwidth manageable as the deployment grows. Preprocessing uses physically meaningful quantities, including a per-segment estimate of shaft frequency derived from the signal's spectral content, so that downstream features remain interpretable in mechanical-engineering terms rather than as abstract latent variables.

Web-based Diagnostic Platform

The engineer/technician-facing layer is a web application that consumes data products generated by edge nodes. It displays the current hour's RMS vibration channels against a rolling weekly baseline for the same pump, shows STFT spectrograms and integrated power spectral density analysis for easy inspection of frequency content, and highlights the hourly model reconstruction-error score with a customizable alarm threshold. The reason for displaying spectrograms alongside the model output is practical: engineers and technicians trust diagnostics they can verify directly. When the model marks an hour as anomalous, the same view presents the spectrogram that influenced the score, enabling faster triage and reducing the cycle from alert to maintenance decision.

ANOMALY DETECTION METHOD

We treat pump anomaly detection as an unsupervised problem. Faulted vibration data from accelerator pumps is rare, unevenly labeled, and biased toward the modes that have failed in the past, so any supervised classifier would generalize poorly to the long tail of degradation modes we actually care about. To address this issue, we deploy unsupervised model ensembles combining machine learning models. Autoencoders sidestep this difficulty by learning a compressed representation of healthy operation and treating reconstruction error as an anomaly score [3-5].

The windowed segment of the pre-processed vibration data described above are input into the ML models; STFT spectrograms generated by the same preprocessing stage are surfaced separately on the diagnostic platform for interpretation. The ML models are trained on a healthy reference period spanning several weeks of operating data, and s occur periodically. The reconstruction error is compared against a baseline established from held-out healthy data; a sustained error above the threshold is flagged as an error. The model architecture, training procedure, performance, and the role of the ML models within the broader anomaly-detection stack are the subject of a separate publication in preparation and are therefore not detailed here.

DEPLOYMENT AND EARLY OPERATION EXPERIENCE

The system is in pilot deployment at APS. Initial instrumentation has focused on two representative loads: the helium pump that supports the bunch lengthening system (BLS), and a set of deionized water pumps for the storage ring (SR). Together these cover the operational regimes of greatest interest: a thermally sensitive cryogenic support pump where any drift has consequences for super conducting radio frequency (SRF) performance, and high-duty water pumps that represent the dominant pump population by count. Hourly model inferences and data visualization are delivered to the diagnostic platform without interruption. We are using early operational feedback to tune alarm thresholds and visualization defaults before scaling the deployment to additional pumps across the facility.

Several engineering lessons are worth highlighting. First, baseline drift occurs regularly: even healthy pumps show seasonal and load-dependent shifts in their spectral signatures, and rolling weekly baselines absorb these shifts far better than a single static reference. Second, mechanical coupling of the sensor matters as much as the sensor itself; consistent mounting torque and location had a greater impact on data quality than incremental improvements to the model. Third, decoupling the inference pipeline from the visualization stack has let us update one without disrupting the other, which has been important for keeping the system live during iteration.

OUTLOOK

The near-term focus is to expand sensor coverage across the APS cooling network and integrate the alarm channel into existing operations workflows so that flagged events trigger a defined response rather than merely updating the web dashboard. In parallel, we are extending the modeling work to achieve sharper separation of fault classes and earlier detection of slow degradation; results from that line of work, together with quantitative detection performance on operational data, will be reported separately.

The broader point is that low-cost MEMS sensors, modest edge compute, and an unsupervised deep model are sufficient to add a useful predictive-maintenance layer on top of existing accelerator infrastructure. Each of the individual ingredients is a mature technology; what has been missing for pumps at light sources is integration and test results from a pilot deployment.

CONCLUSION

We have described a pilot deployment of a pump anomaly-detection system at APS that integrates low-cost MEMS vibration sensors, Linux edge nodes, an unsupervised machine learning model for hourly anomaly scoring, and a web-based diagnostic platform with STFT spectrogram visualization for engineer or technician facing triage. The system is currently running on the BLS helium pump and on selected storage-ring water pumps. Early operational experience supports the design choices behind the system and motivates a facility-wide rollout. A detailed treatment of the model and quantitative results is forthcoming.

ACKNOWLEDGEMENTS

We thank the APS Engineering Support (AES) group for assistance with sensor installation, and members of the AOP group for helpful discussions on the anomaly detection methodology. We also thank Curt Preissner and Kamlesh Suthar from APS for valuable input on sensor use and calibration.

REFERENCES

[1] E. Swetin, M. Kirshenbaum, and C. Putnam, "Cooling Water Systems for Accelerator Components at the Advanced Photon Source," in *Proc. MEDSI'02*, Argonne, IL, USA, Sep. 2002.

[2] R. B. Randall, *Vibration-based condition monitoring: industrial, automotive and aerospace applications*, John Wiley & Sons, 2021. doi:10.1002/9780470977668

[3] J. Zhang, H. Xia, Z. Wang, Y. Zhu, and Y. Fu, "Research on unsupervised condition monitoring method of pump-type machinery in nuclear power plant," *Nucl. Eng. Technol.*, vol. 56, no. 6, pp. 2220–2238, Jun. 2024. doi:10.1016/j.net.2024.01.031

[4] G. Michau, Y. Hu, T. Palmé, and O. Fink, "Feature Learning for Fault Detection in High-Dimensional Condition-Monitoring Signals," in *Proc. the Institution of Mechanical Engineers, Part O: Journal of Risk and Reliability*, vol. 234, no. 1, pp. 104–115, Feb. 2020. doi:10.1177/1748006X19868335

[5] J. P. Edelen and N. M. Cook, "Anomaly Detection in Particle Accelerators using Autoencoders," in *Proc. the 2021 Improving Scientific Software Conference*, Mar. 2021, pp. 5-11. doi:10.26024/p6mv-en77

[6] Y. Sun, "Anomaly Detection by Principal Component Analysis and Autoencoder Approach," in *Proc. IPAC'21*, Campinas, Brazil, May 2021, pp. 1502-1504. doi:10.18429/JACoW-IPAC2021-TUPAB061

[7] W. Blokland *et al.*, "Uncertainty aware anomaly detection to predict errant beam pulses in the Oak Ridge Spallation Neutron Source accelerator," *Phys. Rev. Accel. Beams*, vol. 25, no. 12, p. 122802, Dec. 2022. doi:10.1103/PhysRevAccelBeams.25.122802

[8] M. I. Radaideh *et al.*, "Time series anomaly detection in power electronics signals with recurrent and ConvLSTM autoencoders," *Digital Signal Process.*, vol. 130, p. 103704, Oct. 2022. doi:10.1016/j.dsp.2022.103704

[9] R. Sainju, M. Borland, O. Mohsen, and Y. Sun, "Application of low-cost sensors and deep autoencoders for monitoring water pumps in particle accelerators," in *Proc. NAPAC'25*, Sacramento, California, USA, Aug. 2025, pp. 694-698. doi:10.18429/JACoW-NAPAC2025-WEP006

[10] STEVAL-STWINBX1 — ST MEMS wireless industrial node evaluation kit, STMicroelectronics, https://www.st.com/en/evaluation-tools/steval-stwinbx1.html