INL/EXT-21-61772
Revision 0

Analytics-at-scale of Sensor Data for Digital Monitoring in Nuclear Plants

February 2021

2nd Annual Report

Cody Walker, Vivek Agarwal, Nancy Lybeck
Idaho National Laboratory

Pradeep Ramuhalli
Oak Ridge National Laboratory

Mike Taylor
Electric Power Research Institute

INL is a U.S. Department of Energy National Laboratory
operated by Batelle Energy Alliance, LLC

INL/EXT-21-61772
Revision 0

Analytics-at-scale of Sensor Data for Digital Monitoring of Nuclear Plants

2nd Annual Report

Cody Walker, Vivek Agarwal, Nancy Lybeck
Idaho National Laboratory
Pradeep Ramuhalli
Oak Ridge National Laboratory
Mike Taylor
Electric Power Research Institute

February 2021

Idaho National Laboratory
Idaho Falls, Idaho 83415

http://www.inl.gov

Prepared for the
U.S. Department of Energy
Office of Nuclear Energy
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517

Page intentionally left blank

ABSTRACT

For economic reasons, the nuclear industry is witnessing premature closure of nuclear power plants, despite excellent safety records. Operations and maintenance activities are some of the largest costs in operating legacy light-water plants. By reducing operations and maintenance costs, nuclear energy can become more economically competitive with other energy sources. This can be achieved in part by leveraging machine-learning and artificial intelligence technologies to develop data-driven algorithms to better diagnose potential faults within the system. The improved accuracy of the models can lead to a reduction in unnecessary maintenance, thus reducing costs associated with parts, labor, and unnecessary planned, forced, or extended outages.

To address these challenges, the goal of this project is to perform research and development in the area of digital monitoring (i.e., the application of advanced sensor technologies, particularly wireless sensor technologies, and data-science-based analytic capabilities) to advance online monitoring and predictive maintenance in nuclear plants and improve plant performance (efficiency gain and economic competitiveness).

This report summarizes the Fiscal Year 2020 research progress encompassing the (1) different wireless vibration sensor and data indicators used to assess the health of a plant asset; (2) development of diagnostic models for fault detection; and (3) development of prognostic models for estimating the health of the system up to 7 days ahead.

Page intentionally left blank

CONTENTS

FIGURES

TABLES

Page intentionally left blank

ACRONYMS

BOP	balance-of-plant
BWR	boiling-water reactor
CBP	condensate booster pumps
CBM	condition-based maintenance
CHIP	component health indicator program
CM	corrective maintenance
CLOP	condensate lube oil pump
CMMS	computerized maintenance management system
CP	condensate pumps
SAD	distributed antenna systems
EPRI	Electric Power Research Institute
FWCS	feedwater and condensate system
IEEE	Institute of Electrical and Electronics Engineers
LTE	long-term evolution
ML	machine learning
NPP	nuclear power plant
O&M	operations and maintenance
OLM	online monitoring
PC	primary component
PM	preventive maintenance
PMBD	Preventive Maintenance Basis Database
PMO	preventive maintenance optimization
PWR	pressurized-water reactor
RF	radio frequency
RPV	reactor pressure vessel
SVM	support vector machine
VSN	vibration sensor mode

Page intentionally left blank

Analytics-at-scale of Sensor Data for Digital Monitoring of Nuclear Plants

1. INTRODUCTION

For economic reasons, the nuclear industry is witnessing the premature closure of nuclear power plants (NPPs), despite excellent safety records [1]. Operations and maintenance (O&M) activities are some of the largest costs in operating legacy light-water plants [2]. By reducing O&M costs, nuclear energy can become more economically competitive with other energy sources. Leveraging machine-learning (ML) and artificial intelligence technologies to develop data-driven algorithms to better diagnose potential faults within the system can help decrease O&M costs [3]. The improved accuracy of the models can lead to a reduction in unnecessary maintenance, thus reducing costs associated with parts, labor, and unnecessary planned, forced, or extended outages.

The nuclear industry's excellent safety record is partly due to the amount of preventive maintenance (PM) it performs. Inspections, replacements, refurbishments, vibration monitoring, oil sampling, and other such tasks are all performed at regular intervals [4]. Each task carries an associated cost but reduces the likelihood of certain component failure modes. The frequencies at which these tasks are performed should be reassessed throughout the life of the plant as a means to reduce O&M costs while still maintaining safety.

To address these challenges, the goal of this project is to perform research and development in the area of digital monitoring (i.e., the application of advanced sensor technologies, particularly wireless sensor technologies, and data-science-based analytic capabilities) to advance online monitoring and PM in nuclear plants and improve plant performance (efficiency gain and economic competitiveness). Specific project objectives are:

1) Design a general methodology for a technoeconomic analysis of wireless sensor modalities for use in monitoring equipment condition, especially in balance-of-plant systems in an NPP.
2) Apply data-science-based techniques for decision making and discovery to develop and evaluate integrative algorithms for diagnostic and prognostic estimates of equipment condition using structured and unstructured heterogeneous data distributed across space and time (i.e., analytics-at-scale), including new data from wireless sensors in a nuclear plant.
3) Develop a visualization algorithm to present the right information to the right person in the right format at the right time.
4) Validate the developed approaches and algorithms using independent data from an operating plant.

Critical to addressing these objectives will be the development of appropriate metrics (e.g., mean time to fault detection and confidence bounds) that help quantify the right information at the right time. In addition, there is a need to determine whether advances in sensor technologies (particularly wireless sensor technologies) are able to address current technical gaps in monitoring for nuclear plant equipment.

In the first year of research, the project team developed a general methodology for the technoeconomic analysis of wireless sensor modalities for use in monitoring equipment condition [5, 6].

In the second year of research, the project team focused on developing and evaluating integrated algorithms for diagnostic and prognostic estimation of equipment condition in collaboration with Exelon Generating Station. The feedwater and condensate system was used as the identified balance-of-plant system. Data associated with the system was provided to Idaho National Laboratory and Oak Ridge National Laboratory for analysis. Electric Power Research Institute provided technical guidance and direction on the project. They also provided essential reference documents related to fault conditions observed in the feedwater and condensate system.

This report summarizes the second-year research progress, encompassing (1) different wireless vibration sensor and data indicators used to assess the health of a plant asset; (2) the development of diagnostic models for fault detection; and (3) the development of prognostic models for estimating the health of the system up to 7 days ahead. The report is organized as follows:

- Section 2 discusses wireless sensor modalities that could be used to collect vibration data and rely on different communication protocols
- Section 3 discusses nuclear power fault diagnostics and PM optimization
- Section 4 discusses development of prognostic models using plant asset data
- Section 5 presents a summary and discussion of the FY 2021 research plan.

2. WIRELESS SENSOR MODALITIES TO COLLECT VIBRATION DATA

A milestone report [7] was issued in May 2020 to support the installation of wireless sensors on plant assets; a wireless architecture based on a distributed antenna system (DAS) to support multiple communication types, was required. Idaho National Laboratory (INL) proposed a wireless network deployment strategy for an NPP [1] that would enable an application with low- to high-power needs, low- to high-frequency ranges, and short- to long-range communications, as shown Figure 1. INL has already performed a technoeconomic analysis of the wireless architecture in Figure 1 [5, 6]. The whole network topology is predominantly operated using a DAS long-term evolution (LTE) system or wireless local area network system, since they can enable:

- High bandwidth and data transmission rates with low latency
- Prioritized data transmission, based on the required quality of experience or quality of service
- Most of the wireless technologies to have either a Wi-Fi or a DAS system as their backend network (e.g., a Bluetooth device can connect to Wi-Fi in the back end to upload its data to the internet)
- The system to act as a bridge between end devices or other wireless technologies and the internet or an outside network
- Easy network maintenance by bringing all the networking technologies under one network architecture.

A summarization of wireless vibration sensors and data indicators used to assess the health of a plant asset was performed. Different wireless vibration sensors are considered based on the availability of network compatibility and ease of deployment.

2.1 Vibration Data

Most motors and pumps in an NPP are monitored via periodic vibration monitoring to assess their operating condition based on set of well-defined features. Motors can be categorized into synchronous or asynchronous motors, depending on its application and usage. Also, these motors have different types of bearings. Commonly observed bearing types are roller bearing and sleeve bearings. The vibration measurements collected on these motors and pumps at fixed locations are useful in assessing their conditions. The time series data is converted into frequency domain for assessment. The vibration signal collected at bearing locations have specific frequency characteristics, referred to as fault signatures, that are used to assess the condition of bearings. By using manufacturer data on machine faults, degradation

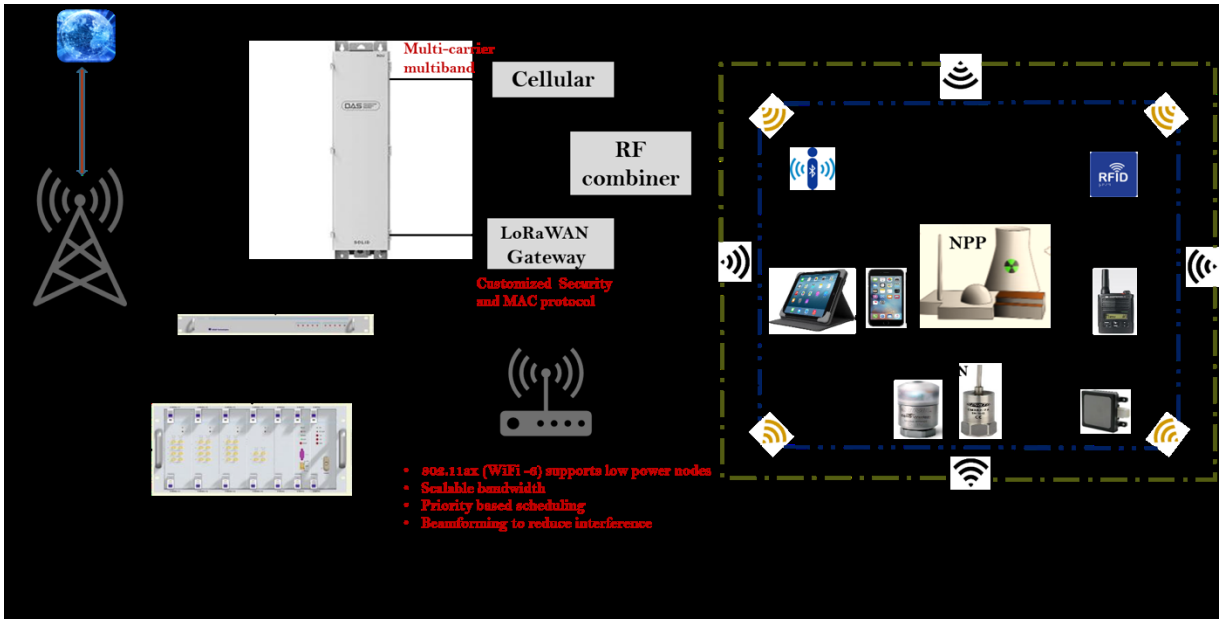

Figure 1. Envisioned multiband wireless technology for nuclear power plant automation [5, 6].

modes can be determined by comparing the frequency spectrums of the measured data and manufacturer data. There are also techniques that can assess the severity of the bearing faults.

2.1.1 Rolling Element Bearing Machines

Rolling element machines can be expected to generate high-frequency vibrations if a bearing fault is developing. The analysis parameter should monitor up to 70 orders of rotational speed and can be expected to provide an early warning of a bearing fault. An acceleration band in the 1–20 kHz range should be used for trending purposes and can be utilized to provide early fault detection.

2.1.2 Sleeve (Journal) Bearing Machines

Sleeve bearing machines are expected to generate fault frequencies at lower multiples of running speeds than a rolling element bearing machine; therefore, it is not generally necessary to monitor as high of a multiple of the shaft speed. Sleeve bearing machines can also generate fault frequencies at subsynchronous frequencies; therefore, a subsynchronous analysis band is needed. Here subsynchronous frequencies refer to the frequencies present below the running speed of the motor and pump.

2.1.3 Gearboxes

Gearboxes can generate fault frequencies at very high multiples of shaft rotation, due to gear mesh frequencies. In order to effectively monitor for gear faults, the analysis parameter should monitor up to 120 orders of running speed with a high resolution to allow the identification of sideband frequencies.

2.2 Plant System Indicator from Vibration Data

Once the raw vibration data is safely stored in the cloud, the data is accessed through an installed sensor cloud server, which provides an interface to visualize the data and evaluate the state of the equipment. One of the steps performed by the plant engineers is to generate plant indicator data from raw vibration data. A single sensor node can trend dozens of different indicators for plant use.

The raw data consists of the asset's acceleration with time. From this time domain data, an acceleration frequency spectrum is calculated by executing a Fast Fourier Transformation routine. It is also possible to convert the acceleration data into velocity signals in both time and frequency domains. Acceleration generally highlights high-frequency vibrations. Velocity tends to be more effective for

evaluating vibration throughout the entire frequency spectrum. Thus, four base-data types are available to be processed into indicators:

- Acceleration time waveform (raw measured signal)
- Acceleration frequency spectrum (calculated)
- Velocity time waveform (calculated)
- Velocity frequency spectrum (calculated).

From these four data types, useful quantitative indicators (mentioned below) can be calculated and trended to characterize machine vibration. Users or plant engineers can then decide which of the available indicators are important to track for their specific applications:

- Peak: Acceleration, Velocity
- Root Mean Square: Acceleration, Velocity
- Spectrum Overall: Acceleration, Velocity.

2.3 Wireless Vibration Sensor Technologies

Vibration-monitoring instrumentation contains accelerometers that sense changes in the amplitude and frequency of dynamic forces that can impair rotating equipment. Identifying degradation at its onset by analyzing vibration measurements allows personnel to identify issues, such as imbalance, looseness, misalignment, or bearing wear in assets prior to significant degradation and failure. This gives the plant more options and more time to respond, allowing for more effective resolutions.

There are wireless vibration sensors that are utilized by plant sites to measure vibration. These wireless vibration sensors could be uniaxial, biaxial, or triaxial. Also, these wireless vibration sensors can have different wireless connectivity characteristics, such as communicating information over a cellular gateway, enterprise Wi-Fi, or 900 MHz. Among different wireless connectivity modes, the performance of cellular and Wi-Fi are comparable, while 900 MHz has a limited capability to transmit measured vibration spectrum or waveform. A basic architecture to support wireless vibration data transmission from sensor to a dashboard is shown in Figure 2.

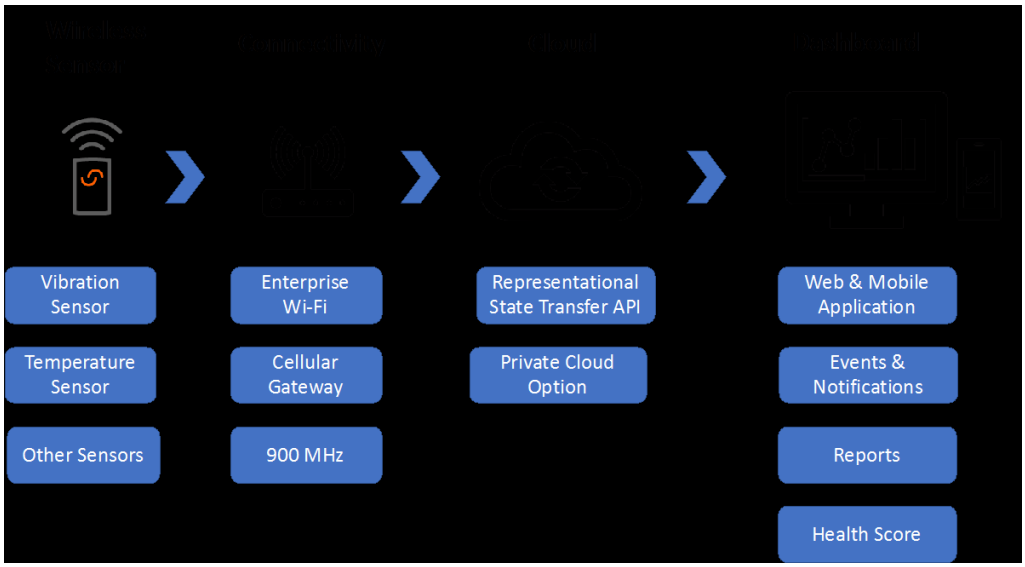

Figure 2. General schematic of wireless data transmission from sensor node to base station for processing, analysis, decision making, and visualization.

When an adverse trend or equipment degradation is detected from the analysis parameters on a dashboard, the Vibration Technology Program Owner must perform an update to the component health indicator program (CHIP) database in a timely manner. The status of each component and overall component health is maintained in the CHIP database. All issues reported in CHIP require the creation of a problem report, as well as a CHIP technology examination, with a notification. The CHIP technology examination and notification document the component condition and recommended corrective actions.

This section discusses different types of wireless vibration sensors compatible with different network communication protocols along with proximity probes. These are some of the wireless vibration sensors that partner plant sites are considering installing on their plant assets to enable online vibration measurement.

2.3.1 Wi-Fi-Based Wireless Vibration Sensor

There are many commercially available Wi-Fi-based wireless vibration sensors. These vibration sensors transmit data over the established enterprise Wi-Fi network from a measurement point to a gateway and from the gateway to the cloud. One of the Wi-Fi-based wireless vibration sensors is Petasense [8]. The Petasense wireless vibration sensor, referred to as a vibration mote and shown in Figure 3, contains three accelerometers and a thermocouple. The vibration mote is battery powered and can be epoxy or stud mounted on plant assets.

Figure 3. Petasense triaxial vibration mote [8].

2.3.2 IEEE 802.15.4 Based Wireless Vibration Sensor

One of the wireless vibration sensors based on IEEE 802.15.4 wireless is the Bently Nevada Ranger Pro, as shown in Figure 4 [9]. This sensor measures velocity, acceleration, and temperature. Ranger Pro wireless vibration sensors are efficient for machines with roller-element bearings, which include, but are not limited to, centrifuges, motors, small reciprocating compressors, and small hydro and steam turbines. The wireless transmission range of these sensors varies depending on environmental obstacles, gateway antenna type, and the orientation of the sensor relative to the gateway antenna.

Figure 4. Bently Nevada Ranger Pro wireless vibration sensor [9].

2.3.3 Cellular Compatible Wireless Vibration Sensor

The KCF Technologies wireless sensor node transmitted X- and Y-direction vibration data, along with temperature data, to the base station over a Dynamic Address RouTing (DART) wireless protocol, which can then transmit the data to the cloud either by using the Wi-Fi network or cellular network [10]. The data stored in the cloud is accessed through KCF SMARTDiagnostics® (SD) machine condition monitoring software [11]. SMARTDiagnostics® provides an interface to visualize the data and the state of the equipment. Each sensor node contains two accelerometers and a thermocouple. The sensor node contains a battery and a transmitter to communicate data to the base station. The vibration sensor node (VSN) is software-configurable and scalable. Hundreds of sensor node points can be accommodated, and each sensor node can be configured to transmit data at a user-selected frequency. Unique indicators derived from the data can be implemented to alert users of potential machine health issues. The dimensions of the SD-VSN-3 are given in Figure 5.

SMARTDiagnostics® uses the 2.4 GHz industrial, scientific, and medical (ISM) Radio Frequency (RF) band for sensor node communication. The 2.4 GHz ISM RF band is able to effectively balance energy efficiency and the transmission range within industrial facilities. Other bands, such as the 915 MHz ISM band, offer an insufficient data rate to transfer large data sets effectively. The 5.8 GHz band has a range is too close for practical application in industrial conditions.

SMARTDiagnostics® makes use of five channels, as shown in Figure 6. The five channels are overlaid on a Wi-Fi spectral usage chart. Wi-Fi typically operates within hannels 1, 6, or 11. If Wi-Fi Channel 6 happens to be used in the plant, SMARTDiagnostics® Channels D and E can be enabled. If Wi-Fi Channels 1 and 11 are active, SMARTdiagnostics® Channels A, B, and C can be used without disruption.

SMARTDiagnostics® uses one or more of the five available channels, depending on the application. Each channel usually communicates with up to 50 sensor nodes. The wireless coverage area is approximately 200 feet from the base station. Within the 200-foot perimeter, 250 sensors can usually communicate with a single collection server in the cloud. Industrial plants can be split into defined areas, with collection servers disseminated to support each area.

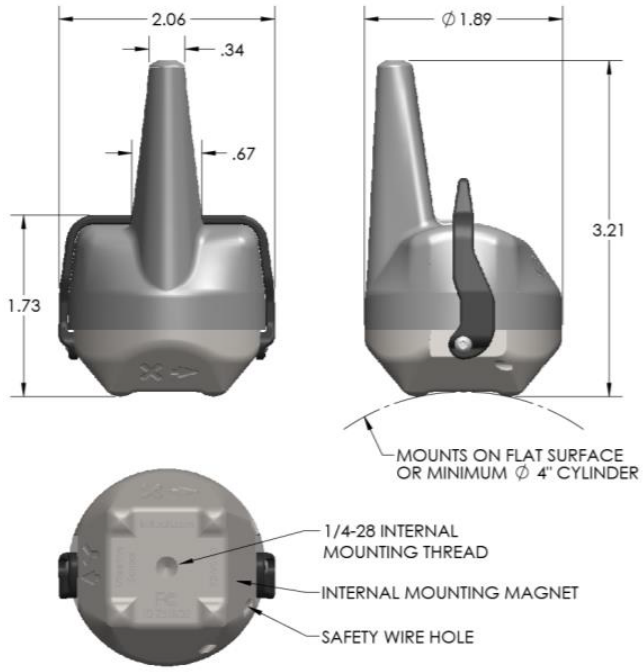

Figure 5. The dimensions of the SD-VSN-3 sensor node [10].

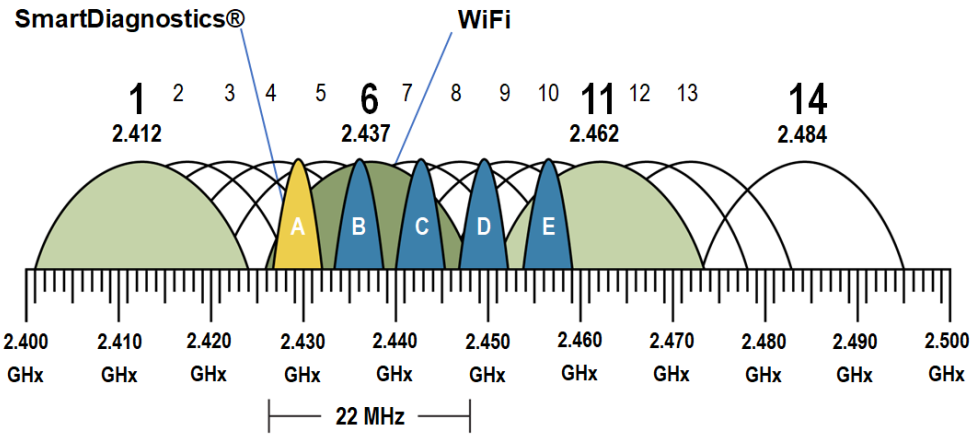

Figure 6. SMARTDiagnostics® uses a modified portion of the 2.4 GHz ISM RF band for the sensor node transmission [11].

2.3.4 Proximity Vibration Probe

Traditionally, plants have relied on proximity probes to measure any vibrational changes in rotating machinery. Proximity probes use the Eddy current principle to measure the distance between the probe tip and the surface to be observed. The proximitor, an electronic device, generates a low power RF signal that is connected to a coil of wire inside the probe tip to the extension cable to measure changes in the field or return signal without physical contact. Figure 7 shows a Bently Nevada proximity sensor system [12]. A proximity sensor system is commonly used to measure vibrations in rotating machines with sleeve bearings. Proximity probes are available in many different lengths and diameters.

Figure 7. Bently Nevada proximity sensor system [12].

3. NUCLEAR POWER FAULT DIAGNOSTICS AND PREVENTIVE MAINTENANCE OPTIMIZATION

A milestone report issued in January 2021 [13] documented achievement of the objective to create a fault diagnostic model via a two-prong approach: condition-based monitoring (CBM) and PM optimization (PMO). This report first investigated CBM techniques in order to identify potential faults and degradation within the feedwater and condensate system (FWCS) condensate pumps (CPs), condensate booster pumps (CBPs), and associated motors. This report focused on utilizing heterogenous data in combination with information from the computerized maintenance management system (CMMS) in order to identify potential condition indicators, which were then used in conjunction with ML techniques to seek out other faults within the system. By accurately predicting degradation, the appropriate PM can be scheduled in order to avoid unplanned downtimes and asset failures. Actual plant data from a boiling-water reactor (BWR) FWCS over a large period of time was used in this CBM analysis. Next, this report examined the application of PMO to update PM frequencies, based on component histories and the opinion of subject matter experts. This research focused on how to determine whether a component is a PMO candidate, along with the steps to perform the PMO so that labor costs can be minimized by reducing PM frequencies in an intelligent manner. CBM and PM work orders from a pressurized-water reactor (PWR) over a 5-year period were used in the PMO analysis.

3.1 Data Description

Different amounts of data were available for the BWR and PWR systems. For the BWR system, sensor data and CBM records were provided for a 5-year period. No PM records were available, so a final PMO cannot be made. However, an assessment of whether PMO is applicable can be made using the data and records available. For the PWR system, a year's worth of sensor data was available for two units, while CM and PM records were available for a 5-year period. This combination is better suited for PMO, since all the maintenance tasks are available. This dataset may be less suitable for fault diagnostics, depending on the signal of interest. Some temperature signals experience a clear seasonal trend: colder in winter, warmer in summer.

The available sensor data from both systems are similar and include variables such as generator gross load, average feedwater flows, and data on temperatures and pressures for the condensers, CPs, CBPs, and associated motors for multiple pump trains. Each dataset consists of unlabeled data and is sampled hourly. There was no indication as to whether any portion of the data corresponded with equipment failure. The available datasets each encompassed steady-state operation, shutdowns for refueling, and derates of varying sizes. There was insufficient information to determine the cause of each derate.

Because of the varying states of operation, the data was divided into four categories: steady state, refueling, derate, and trip. The steady state covered all instances when the reactor was above 90% of the nominal full power. The power still fluctuates within this category; it is still covered within the broadly labeled steady state. Refueling covers the period of time between the initial ramp down through the refueling outage. Derates contain all observations in which the reactor is operating between 5 and 90%

nominal full power. Trips are similar to derate, but, instead of the gross load being reduced, the reactor is shut down.

The data was cleaned before being processed in the fault diagnostic or prognostic models. The cleaning process addresses issues with missing data, outliers, seasonal variations, and the separation of data into training and testing sets. Missing data was primarily due to the component being offline. If the component was online but the data were not recorded due to an error in the sensor or data archiving process, the value was linearly interpolated using its neighboring values. Potential outliers were selected based on values being at least four standard deviations away from the average for the steady-state operation data. Once an outlier was found, it was replaced using a median filter with a window size of 51, as seen in Figure 8. Temperature data containing seasonal trends were median-filtered with a window size of 700 to remove seasonal variations and long trains of outliers.

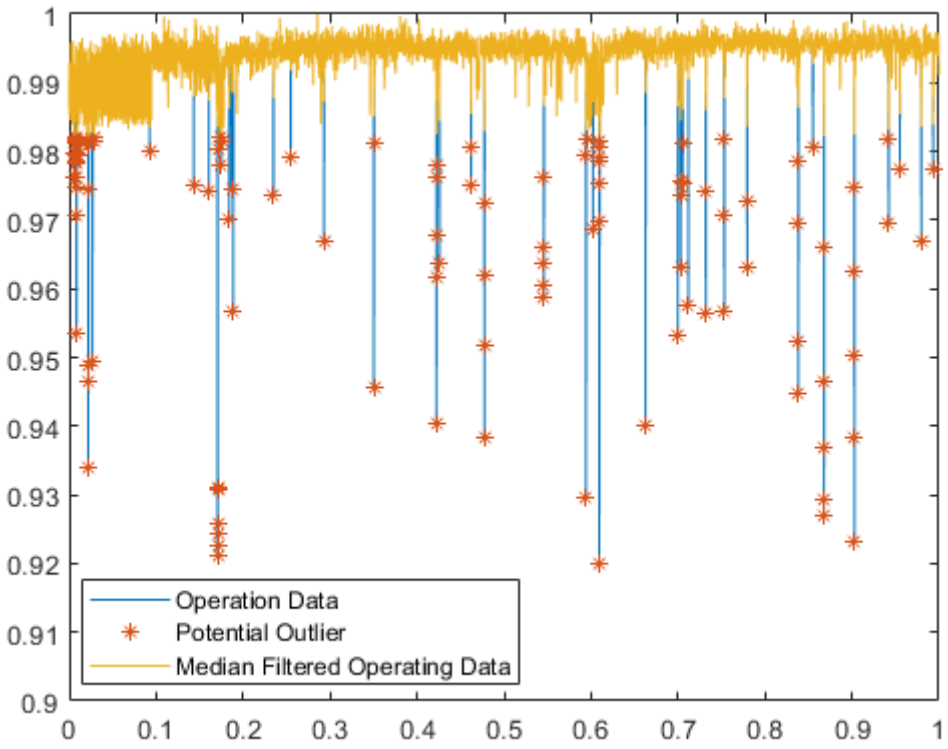

Figure 8. Example of outliers identified in the average total feedwater flow measurement. The data range has been normalized so that the maximum value of the variable is 1 and that the normalized horizontal axis represents the fraction of plant operation time.

3.2 Fault Diagnostic Model

This report analyzes fault diagnostics from two perspectives: a data-driven perspective and a CMMS perspective. In the data-driven perspective, the sensor data is analyzed for faults, anomalies, or trends that may indicate degradation or reduced performance. The CMMS is then used to provide the ground truth to verify that a fault occurred, based on any CM work orders remedying the issue. Some changes seen within the FWCS sensor data are not due to faults but rather changes in operating conditions related to other systems. Diagnosing this behavior is difficult, as sensor data from these other systems were unavailable for this report. The fault diagnostics process from the CMMS perspective begins by analyzing the CMMS for CM work orders related to failures and fixes. The sensor data collected prior to the work order being issued are analyzed for signs of incipient degradation or developing trends over a period of time. It should be noted that not all failure modes can be detected using the sensor data available. Types of available sensor data include gross load, bearing temperatures, pressures, flow, and motor current.

Failures concerning seals, shafts, impellers, and oil require some combination of inspection, vibration, and oil analyses. NPPs perform these tasks at regular intervals, but such data were unavailable for this report. Instead, this report focuses on the process data acquired from in-situ sensors.

A temperature condition indicator measuring the difference between the component's current temperature and seasonal average was created and can be seen in Figure 9. Positive values indicate that the component is running hotter than the seasonal average; negative values indicate that the component is running cooler. Positive values above a 6°F temperature difference were labeled as faulted. This cutoff was chosen based on the component's history before the outage. The faulted section can be seen in the red highlighted area in Figure 9.

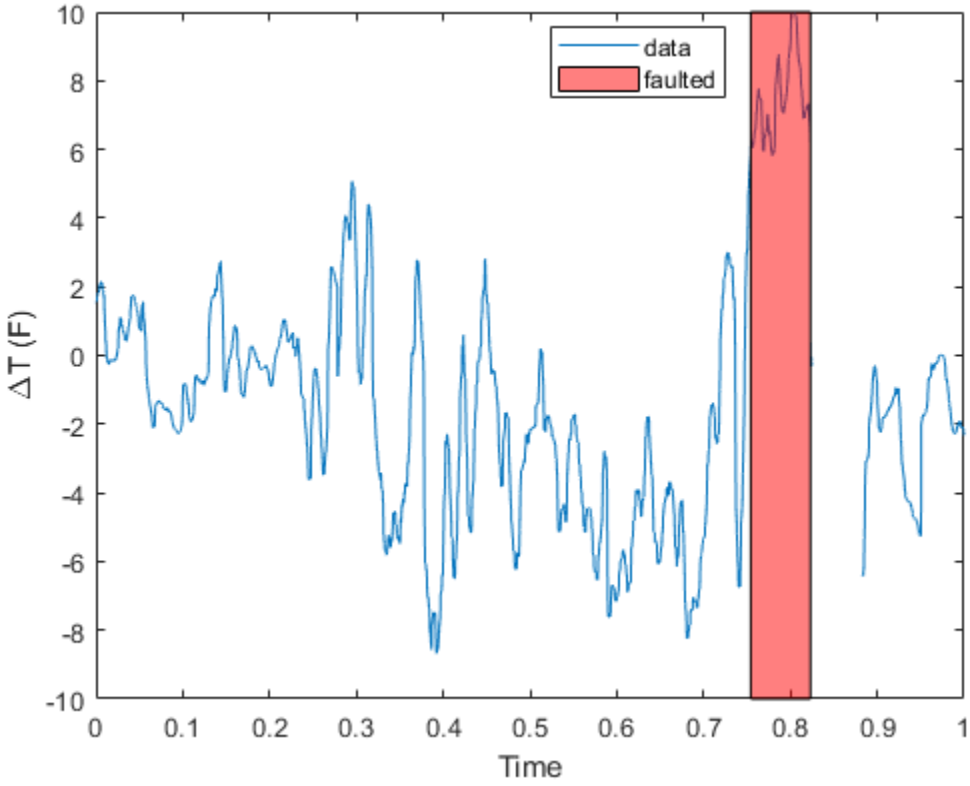

Figure 9. The average seasonal component temperature was subtracted from the current component temperature. Positive values indicate that the component was hotter than the seasonal average. The section of training data highlighted in red is labeled as faulted.

The temperature condition indicator was combined with other process variables (e.g., gross load, flow, and pressure) to observe the total system response to this potential fault. The high dimensionality of this combination was reduced using principal component (PC) analysis. With just three PCs, 87% of the information within the dataset was represented. Separation between the faulted and healthy data may not be due to the existence of a fault but rather the ramp down for a planned outage. The temperature difference between the expected and actual temperature values changes in relation to the plant's other process variables. This nonstationary behavior makes it difficult to determine the cause of the separation.

SVM uses the fault labels to search for similar faults in other years. First, to capture the differences between the healthy and faulted data, the SVM was trained on the data presented in Figure 2. Faulted data made up 4.1% of the total training data. The SVM was trained using a linear kernel function, resulting in 242 support vectors. The SVM was then tested on the training data to double-check that it could correctly label that data. On the training data, it proved 99% accurate. The confusion matrix for the training data showed that it labeled 305 out of 360 faulted data points correctly, while labeling 8,371 out of 8,400

healthy data points correctly. The SVM was then used to search for similar faults from other years. The test data was normalized and transformed into the same PC space as the training data. The SVM then labeled the test data as either faulted or healthy, as shown in Figure 10.

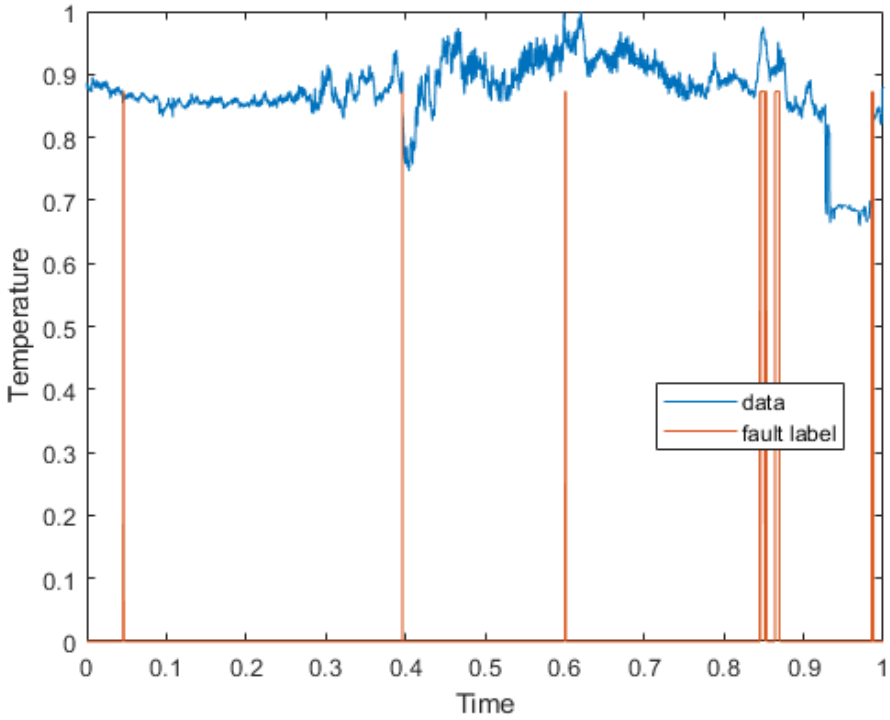

Figure 10. The SVM was used to predict faults in other years. The red spikes indicate either potential faults or inaccurate labels.

The red spikes throughout Figure 10 show each data point that the SVM labeled as faulted. This test year is interesting because no repairs were made to the CP of interest, and a similar de-ramp into an outage was performed. Since no repairs were made to this pump, all data throughout the test year should be labeled as healthy. Under this assumption, the accuracy of the predictions seen in Figure 3 was 81.4%. The accuracy in predicting other years reached as high as 99.7% for years involving only healthy, steady-state operation. This high accuracy can seem misleading, but the system was operating under healthy, steady-state conditions over that entire year, so no faults were expected to be identified.

Fault diagnostics were then analyzed from the CMMS perspective by reviewing the CM work orders related to failures and fixes. CM work orders generated outside the regularly scheduled refueling outages were given higher priority in this review, since they can immediately affect component uptimes and plant output. One such fault was corroborated using the CMMS. An inboard oil leak was discovered in one of the CPs. This leak was due to a bearing's outer edge riding against the coupling side of the housing. The CP was taken offline during plant operation to fix the bearing's leak, causing a 27.5% derate in the plant's output over a 9-hour period. This leak was most likely discovered during a routine walkthrough or via route-based vibration monitoring. The leak was not significant enough to impact those parameters that were being monitored and to which we had access. These parameters included gross load, inboard/outboard bearing temperatures, flow, and pressure. This shows that complete situational awareness is unavailable via the current capabilities of the embedded sensors, causing inspections to remain necessary. A potential way of diagnosing this type of leakage is through the position of the temperature control valve. The temperature control valve position may vary, as the leakage results in less volume to control. Alternatively, the bearing leakage could potentially be discovered via in-situ vibration monitoring. Bearing vibrations may be affected by the size and position of the leak or by how the bearing's outer edge rode against the housing. In such cases, the change in vibration would be detected in

a timelier manner through continuous online monitoring.

3.3 Preventive Maintenance Recommendations

PM is performed at predefined intervals, regardless of the component's current condition. The frequencies at which PM tasks are performed should be reviewed and updated as more operating experience is acquired. This creates a feedback loop in which maintenance frequency impacts component performance—in turn, affecting maintenance frequency. To speed up this feedback loop, current PM frequencies for healthy components are compared against the recommendations in the EPRI PMBD, as seen in Table 1. By reviewing a component's operating history and current PM frequency, a justification can be made for suggesting that the frequency be extended.

Table 1. Example of current PM frequencies and recommendations.

Component	PM Task	Current PM Frequency	EPRI PMBD	Recommendation
CP and CBP	Refurbishment	8 years	As required	Good candidate for frequency extension
	Vibration Monitoring	3 months	3 months	Keep
	Oil Analysis	6 months	6 months	Keep
Motor	Vibration Analysis	3 months	3 months	Keep
	Fan Cleaning	6 months	2 months	Keep
	Oil Analysis	6 months	6 months or 1 year	Good candidate for frequency extension
	Electrical Testing/ Inspection	5 years	4 years	Keep

Over the 5-year period, a total of 705 CBM activities were undertaken between the two PWR systems with regards to the FWCS. These activities vary in location and severity, ranging from simple inspections or alignments to bearing failures and equipment replacements. This report focused specifically on mechanical issues stemming from the CPs, CBPs, and their respective motors. The CBM work orders were reviewed first, in order to ascertain the historical health of the component.

In Unit 1, the CBPs primarily experienced issues concerning mechanical seal leaks. CBP 11 showed three instances of leakage or seal problems. The CPs experienced issues in regard to their connections to the condensate lube oil pump (CLOP). Unit 2 experienced fewer leak issues than Unit 1 but experienced some bearing, sensor, and CLOP issues. The records of the CP and CBP motors for both Units 1 and 2 were nearly impeccable. The only recorded issues were dirty filters in the Unit 1 motors in 2019. That sort of record can open up opportunities to extend PM frequencies for cost reduction purposes.

A summary of PM tasks and the frequencies at which these tasks are performed for each component are given in Table 1. Next to each component's PM frequency is the EPRI-PMBD-recommended frequency based on subject-matter-expert opinion and component reliability data. Among the seven listed PM tasks, four have current frequencies that differ with those recommended by the EPRI PMBD. Based on CM history and the EPRI-PMBD-recommended frequencies, two tasks were deemed

good candidates for PM frequency extensions: pump refurbishment and motor oil analysis.

A "refurbishment" is defined as replacing a component with a spare and then sending it off to be overhauled and repaired. Refurbishments are conducted every eight years during a plant outage, thus avoiding any unnecessary downtimes. Overall, the CPs and CBPs had a positive CM work history. With their EPRI-PMBD-recommended PM frequency of "as required," CP and CBP component refurbishment are a good candidate for PM frequency extension, though the final decision should not be based solely on component performance. An appropriate risk assessment and review of available sensors must be undertaken to determine the health of the pumps. To aid in determining when the pumps should be refurbished, it is recommended that the vibration monitoring and oil analysis continue at their current frequencies. These tasks can be used to monitor component health while trending any observable degradation. Additionally, routine-based vibration monitoring can be replaced with online vibration monitoring (where applicable) in order to further aid in this process. The decision of when to refurbish is a complex one, requiring more information than available for this report.

The motors driving the CPs and CBPs have experienced minimal problems over the last 5 years. The one instance noted in the CM work orders was due to dirty filters in the Unit 1 CP motors. Since the overall health of the components has been good over a significant period of time, PM frequency extensions should be considered in order to reduce the overall cost of maintenance. Since the only issue that the motors faced were dirty filters, and the current PM frequency is already longer than that recommended by the EPRI PMBD, the fan-cleaning frequency should not be adjusted. There is a potential opportunity to extend the frequencies of vibration and oil analyses for the motors. Table 1 lists the EPRI-PMBD-recommended frequencies as being 3 months for the vibration analysis and 6 months or a year for the oil analysis, depending on whether the motor is classified as critical or noncritical, respectively. Either of these frequencies could potentially be extended, perhaps even both. Furthermore, the vibration-monitoring PM task could potentially be replaced by continuous online monitoring, which is why Table 1 suggests that oil analysis to be the preferred candidate for PM frequency extension. In the PWR system, all four CPs must be running for the system to operate at 100% capacity. The motors that drive the CPs are critical components. The CBPs are operated in a two-out-of-three fashion for 100% capacity. Their operation is alternated regularly in order to evenly spread out the accumulated wear. In this case, a backup CBP is present in case one fails to operate. Even with a potential backup, the CBP motors may still be considered critical, and a 6-month oil analysis may be advisable. With a risk analysis and more information about the oil analysis contents, the PM frequency could likely be extended to 9 months or even a year.

4. DEVELOPMENT OF PROGNOSTIC MODELS USING PLANT ASSET DATA

A milestone report was issued in September 2020 [14] documented achievement of the objective to create a prognostic model to estimate predictions of future measurements of the FWCS. This milestone used three different ML methods, long short-term memory (LSTM) networks, SVMs, and a nonlinear autoregressive neural network (NAR) to estimate three different forecast horizons: 1 hour, 1 day, and 1 week. The predictions were compared and contrasted to the BWR plant data over a 5-year period. This milestone discussed the various aspects of data processing and model development that are likely to influence the performance of the prognostic models. This milestone covers problem formulation, parameter selection, model selection, data preprocessing, model development, and model evaluation.

4.1 Data Cleaning and Processing

For this milestone, the objective of the model was to predict plant conditions at one hour, one day (24 hours), and one week (168 hours) into the future. Each of the three model types was used as a way of comparing ML prognostic models (LSTM and NAR) against a statistical prognostic model (SVR). The available time series data from the process variables was cleaned (same as in the previous milestone) and the data from steady-state operations were separated from data corresponding to ramp-downs and derates. The available steady-state data were further partitioned into two groups: a training data set containing the

first 80% of the steady-state data and a test data set containing the remaining 20%. The models were each built using the training data and were evaluated against the test data. Model performance was measured using prediction accuracy and computational complexity associated with deriving the model from the training data.

Each of the three models was used to estimate the prediction performance based on a subset of process variable measurements corresponding to the total feedwater flow and temperature measurements. To ensure a reasonable comparison between the modeling approaches, a subset of process variable measurements (feedwater flow, feedwater temperature, and feedpump/condensate boost pump operating characteristics) was used to build the models. This subset of data showed a high correlation with the generated load from the plant and was directly related to the equipment of interest for degradation detection. Model development and performance evaluation were performed using one of two methods. In the first method, the normalized data were used directly to build and evaluate the models. In the second method, principal component analysis (PCA) [15] was used to identify the largest principal components from the data. These constitute the components with the greatest explanatory power, with the remaining principal components more likely to be related to noise in the data. These data were used in the model development and evaluation. It is expected that the use of PCA will reduce the model complexity by reducing the size of the models and will enhance the robustness of the models by focusing on features that are more likely to be related to the system states of interest. A further benefit of the PCA approach is the decorrelation of the measurements by projecting the data along orthogonal principal axes; this is expected to improve the prediction accuracy.

4.2 Model Results

The prognostic model performance, especially for the ML models, is expected to be sensitive to the model structure, parameters, and data used. To verify this, the LSTM and NAR models were trained with differing numbers of hidden layer neurons; performance was tracked through the time required to train the models (on a dual-core CPU system) and the root mean square error (RMSE) in the prediction. The RMSE values were estimated using the test data set only. The SVR models were used as a baseline for comparison of the other neural network-based models. In most cases, the models were applied to predict the future values of one or more of the measurements. In a few instances, however, the models were applied to predict the generated output of the plant from the measurements.

The initial analysis focused on predicting both plant power output and other sensor data over different prediction time horizons. Different model configurations and training data were used with the NAR model, yielding the results provided in Table 2. The highlighted rows in the table indicate the case in which the data split was 80/20. For comparison, results are also given (rows without highlighting) from when the training data consisted of all data from refueling cycles with similar plant output on average, with data from other refueling cycles used as test data to evaluate prediction performance.

The number of inputs indicates the number of principal components selected for the NAR model. Table 2 shows that, in general, too many nodes can yield just as poor generalization as too few nodes. Increasing the number of inputs does not always increase the performance of the model. By using fewer PCs, the model sees less noise that is uncorrelated from the signals of interest. The results also indicate a dramatic reduction in the prediction error when using an 80-20 data split (highlighted rows). An examination of the data indicated a small but measurable and consistent change in the generated output between the refueling cycles corresponding to the test and training data. The result of this difference (likely due to an increase in the plant generation capacity in this cycle) was a systematic bias in the predictions. This bias was alleviated by introducing an 80-20 split in the training data, so that the training data consisted of parts from all fueling cycles. This initial result also highlighted the need to ensure the training data is representative when using prognostic models that learn from the data.

Table 2. Summary of the NAR model performance under different combinations of nodes, delays, inputs, and training data. RMSE was averaged for all signals estimated.

# of nodes	# of delays	# of inputs	RMSE

10	10	25	2.068
10	2	25	2.110
25	2	25	1.694
35	2	25	0.881
50	2	25	1.759
35	2	78	2.310
35	2	16	0.660
35	2	13	0.460
35	2	10	0.404
35	2	16	1.711
10	2	6	0.167
12	2	6	0.087
16	2	6	0.126

The left plot in Figure 11 shows an example of the prediction results for the 12-node, 2-delay, 6-input NAR model using an 80/20 training/test data split, for three prediction horizons. The figure presents the normalized load (power generated) against time (normalized as a fraction of plant operation time). Separate NAR models were used to predict one hour, one day, and one week ahead. The right plot in Figure 11 shows similar results for the SVR model. This SVR had a linear kernel function and used default settings, resulting in 857, 692, and 146 support vectors, respectively. Parameter optimization will be necessary to optimize the number of support vectors and enhance the generalization ability of the SVR. For the LSTM models, the same 80/20 split was used. The results summarized in Figure 12 (varying prediction time horizons for LSTM models with different numbers of hidden nodes) appear to show a decrease in prediction accuracy with an increasing prediction time horizon. However, the data indicate variation in this trend, with a potential increase in accuracy as the horizon increases to about a day before decreasing again. A comparison against SVR models using linear and Gaussian kernels is also shown in Figure 12.

In the results presented in Figure 11 and Figure 12, the apparent increase in accuracy at the one-day prediction horizon as compared to the one-hour and one-week predictions was unexpected. This is likely a result of the data used in the training and testing and not a general finding on prediction accuracy as a function of prediction time horizon. Indeed, we would expect that the prediction accuracy would be better when attempting to predict closer in time. The specific cause of the unexpected trend observed has not been determined, but likely causes include network size (for LSTM) and a fortuitous periodicity for small derates in which the 12- to 24-hour window helps the models avoid having to predict many derates using steady-state data and vice versa. This latter potential cause is being investigated further to determine if this is an artifact of the sampling rates used to archive and provide the data. It is worth noting that with all models, performance degradation was seen if the output was routed back to the input to enable multiple time step prediction (i.e., using the predicted output at a time step to predict the output at the next time step), as a result of prediction error accumulating over time. Such a scenario might occur when a single model is operating in real time to obtain multiple time horizon predictions. In contrast, the use of a separate model specifically trained to predict one day or one week in advance appears to provide better results.

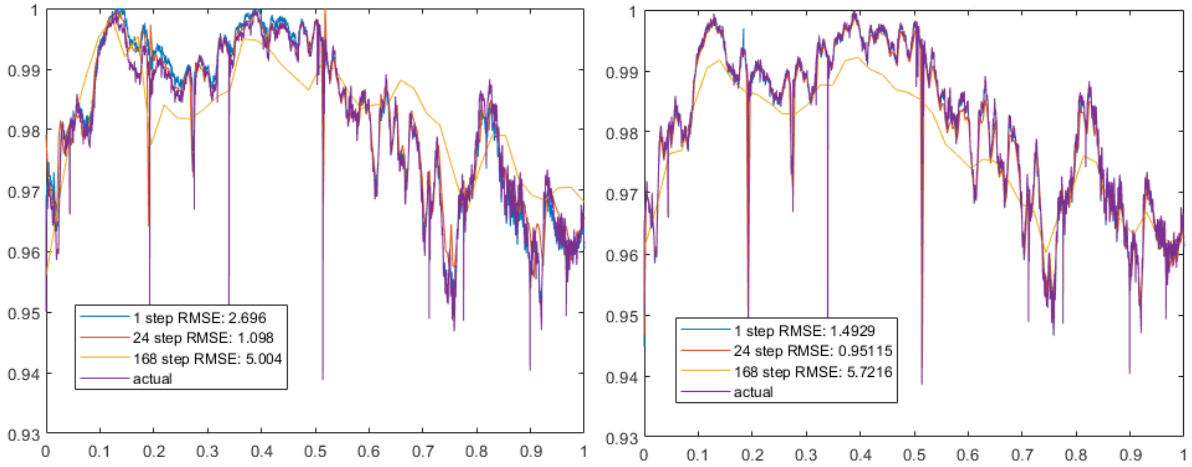

Figure 11. Summary of predictions for the test set (steady-state operations) from the (left) NAR model with 12 nodes and (right) SVR model. Each plot shows the prediction performance for three different prediction time horizons. The horizontal axis is normalized as a fraction of time.

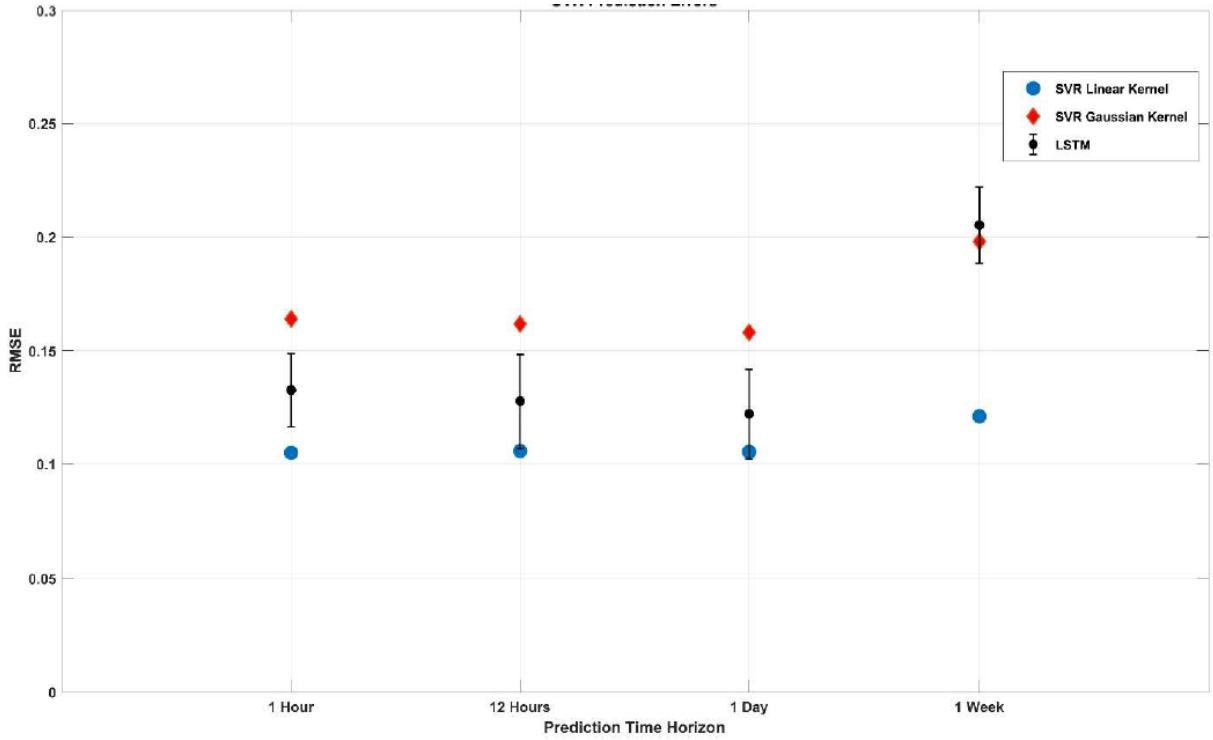

Figure 12. Comparison of RMSE of predictions for the test set (steady-state operations) from LSTM and SVR models as a function of the prediction horizon. The variation in RMSE for the LSTM models are due to models with different numbers of hidden layer nodes.

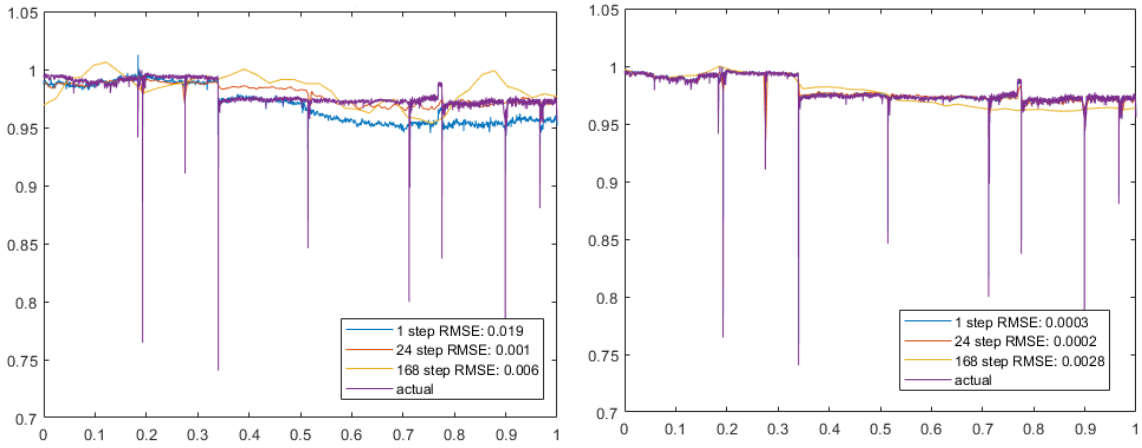

Figure 13. Summary of flow prediction for three different prediction horizons for the test set (steady-state operations) from (left) 12-node NAR model and (right) SVR model, as a function of a million gallons per minute for a fraction of the plant operation time.

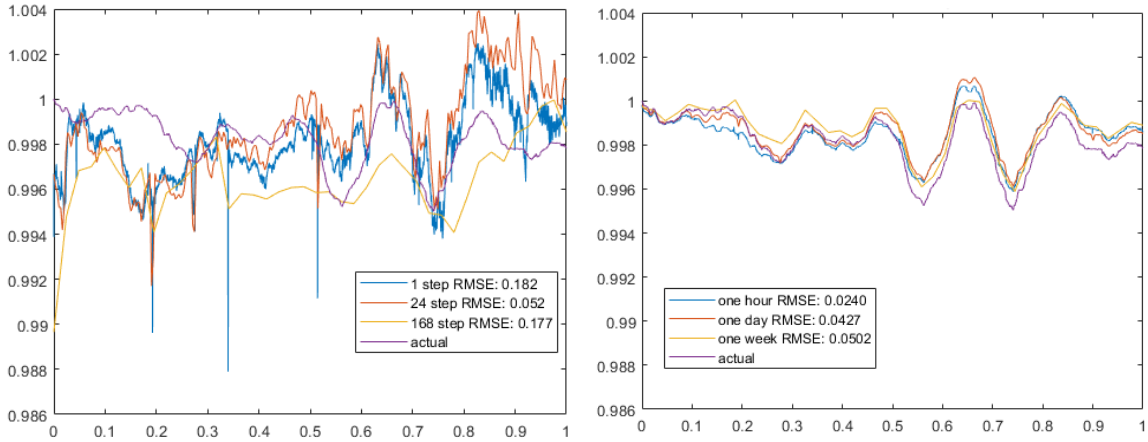

Figure 14. Summary of predictions for three different prediction horizons for the reactor feedwater pump temperature (steady-state operations) from (left) NAR model and (right) SVR for a fraction of the plant operation time.

The models were also used to predict the flow in a single reactor loop for three different prediction horizons. In Figure 13, the NAR prediction (left) is compared against an SVR model (5,903; 411; and 89 support vectors for the three prediction horizons, respectively, shown on the right). Again, the one-day prediction appears to outperform both the one-hour and one-week predictions. However, the NAR model in Figure 13 shows that a consistent bias developed halfway through the test. The cause of this deviation is currently unknown but is being investigated. The SVR outperformed the NAR model by roughly an order of magnitude for each prediction horizon. In Figure 14, similar results were seen when comparing the prediction performance for the reactor feedwater pump temperature of NAR and SVR models (19, 20, and 47 support vectors, for the three prediction horizons for SVR). A similar performance improvement is seen when comparing the LSTM and SVR models (Figure 15); the performance improvement is particularly significant at longer prediction time horizons.

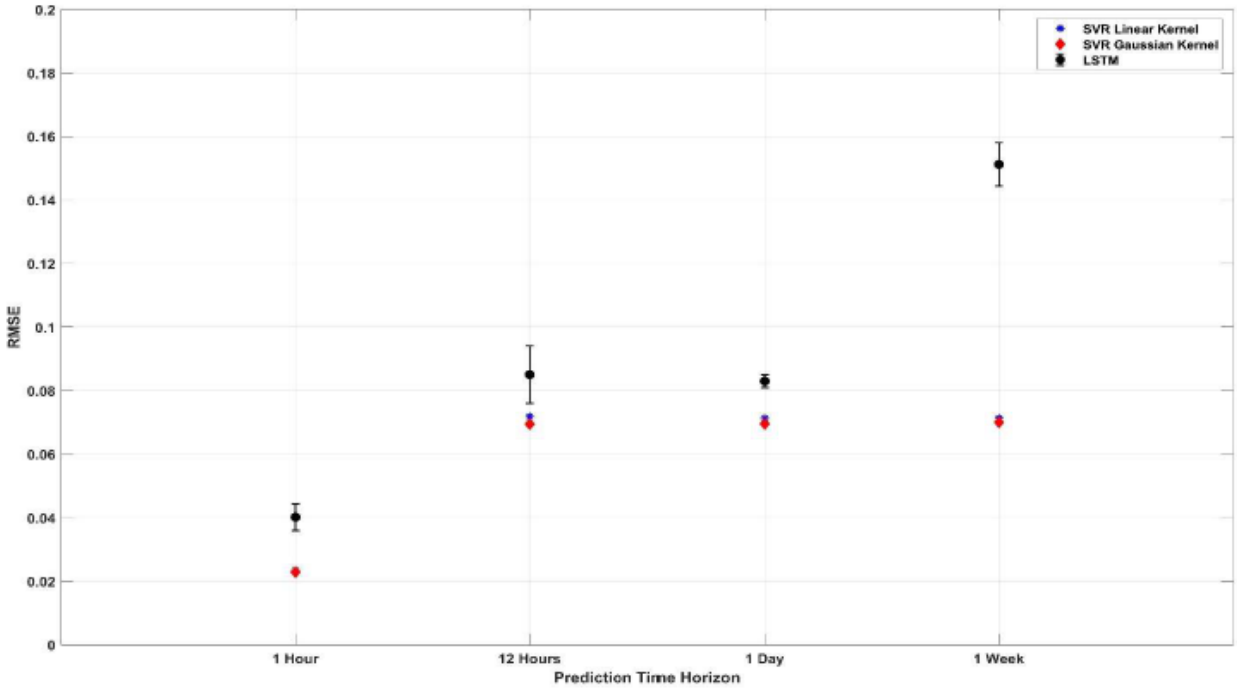

Figure 15. Summary of average flow predictions for the test set (steady-state operations) from LSTM and SVR models as a function of the prediction time horizons.

While the apparent performance advantage for the SVR could be due to the inherent advantages of the model itself, it is more likely that the performance gain is due to the suboptimal model structure for the NAR and LSTM models. The result appears to indicate a need for model structure optimization for LSTM and NAR models and likely for all machine-learned models where model structure selection is a key part of model development. The performance variation of LSTM as a function of the number of hidden layer nodes was seen to vary, as seen in Figure 16 and Figure 17. Variations were also largely due to variation in accuracy at different prediction time horizons. As can be seen from these examples, the prediction accuracy is a complex function of the hidden layer size (larger hidden layers are not always better) and the prediction time horizon. Whereas the hidden layer size (number of LSTM cells) is related to the amount of history from which the LSTM network learns, often there is a limit beyond which increasing the size of the hidden layer brings diminishing returns. Unfortunately, there appears to be no simple way to identify this limit, nor is there a monotonic relationship between the number of hidden layer nodes and the performance of the network before this limit. There is also evidence in the literature that LSTMs do not scale well with data size, indicating that the use of LSTMs as prognostic models for predicting over a very long time series may be challenging. Similar model structure optimization can be performed for SVR by modifying the underlying kernel functions and the kernel bandwidth as well as adjusting other algorithm parameters.

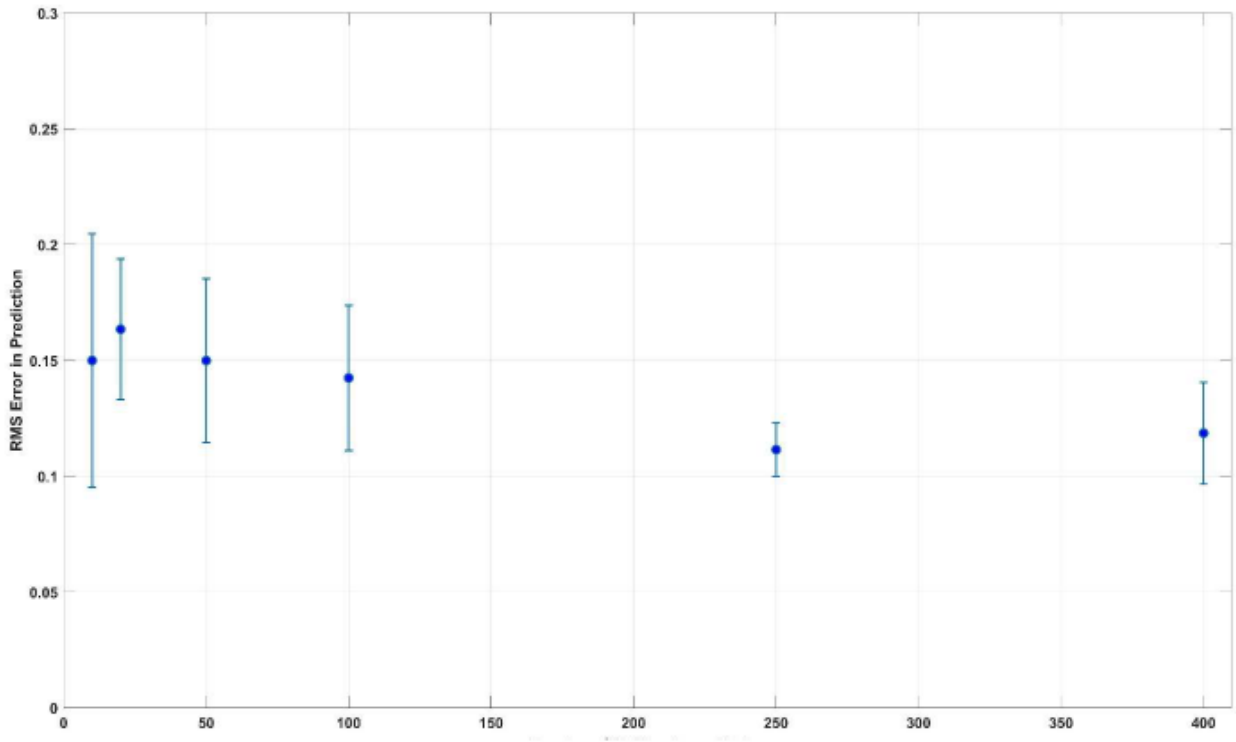

Figure 16. Performance variation of LSTM with hidden layer size. The LSTM was trained to predict the gross load under steady-state conditions using previous measurements of process variables (flow, temperature, pressure). The variation in RMSE for each hidden layer size is due to the different prediction time horizons; a higher RMSE usually correlates with a longer prediction horizon in this instance.

While the assessment was not comprehensive, as it did not examine the effect of each parameter systematically, it did explore a large range within the parameter space, and the findings may be generalizable as a result. The results indicate that, as expected, the time required to train the models increases as the model complexity increases. However, the initialization of the LSTM and NAR parameters was also observed to influence the training time; initialization closer to an optimal solution tends to result in faster convergence. Also obvious from the training process was that the actual training time depends on the data and the algorithm convergence parameters: loss function value at which to stop training, maximum number of epochs for training, mini-batch size, sequence length, etc.

The model performance was found to be sensitive to the normalization of the inputs as well as to whether features—PCA in this instance—were used or not. However, the performance gains with PCA as a preprocessing tool were inconsistent across model types. It is also worth noting that there are several combinations possible for the inputs and outputs to the models, and, to a great extent, the proper choice of which sensors to use as inputs (predictors) tended to be a matter of trial and error given this data set. Alternative metrics, such as the Akaike Information Criterion or Bayesian Information Criterion, are necessary to quantify the information gain from a specific input (predictor) and objectively determine the subset of predictors that should be selected for prognostics.

A key finding here for degradation prognostics is the need for data from the system under monitoring, as well as other associated data from the plant process that may impact performance of the system being monitored. A related issue is the selection of appropriate prognostic models. While ML models appear to be quite capable of prognosing the quantity of interest, relevant labeled data must be available for training the models. Furthermore, the model structure (number of nodes, neural network connectivity, etc.) does not appear to have a monotonic relationship with the prediction accuracy, although there appears to be a monotonically increasing relationship between network size and computational complexity, as in the time

to train the network to reach an acceptable training error and the associated memory requirements. As such, we expect the use of these models to require some trial and error before an acceptable model structure and performance can be achieved.

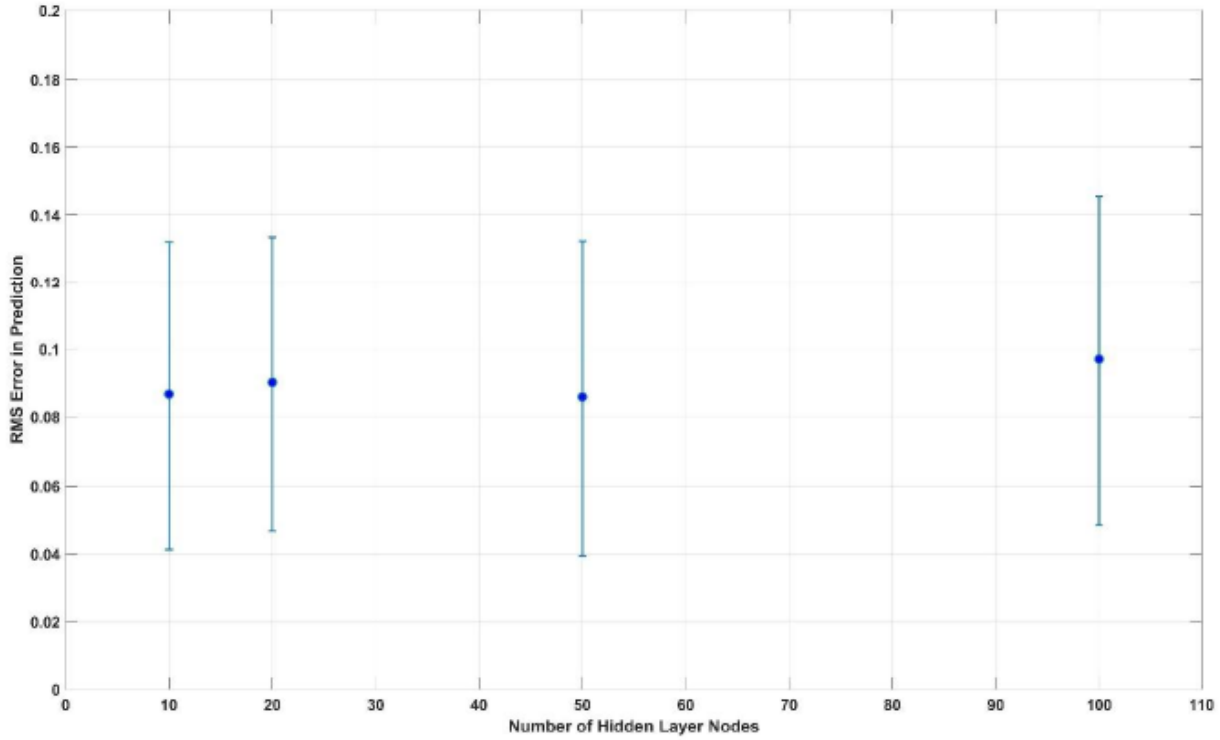

Figure 17. Performance variation of LSTM with hidden layer size. The LSTM was trained to predict the average flow under steady-state conditions using previous measurements of total flow. The variation in RMSE for each hidden layer size is due to the different prediction time horizons: higher RMSE usually correlates with a longer prediction horizon in this instance.

5. SUMMARY AND PATH FORWARD

This report summarized some of the different wireless vibration sensors capable of transmitting data over different communication networks being examined for installation in partner plant assets. We also present an application of vibration data and indicators to assess the health of a plant asset. The report presented an application of SVM, in combination with feature extraction techniques, to search for fault diagnosis. Results indicate that not all faults can be identified via the current suite of embedded sensors. Additional sensors or inspections would be required to locate these elusive faults. This report also showed the analytical steps necessary for PMO, from component health determinations to recommendations for action. This type of analysis can be extended to components within other systems, as long as their performance histories and PM frequencies are known. Finally, the report presented a series of studies using three types of exemplar models and data from an operating plant indicates that the models can learn relatively quickly and may be used to predict the process measurement of interest over different time horizons. The prediction accuracy is dependent on the model structure and other factors, including the data used in the training process, with the accuracy degrading as the time horizon increases. Results indicate that the prediction accuracy tends to be highest when working with data from steady-state conditions; the use of data that includes transients challenges the ability of the model to predict the behavior during and sometimes after the transient as occurred.

As part of the research in FY 2021, the project team will work towards these objectives:

1) Develop a visualization algorithm to present the right information to the right person in the right format at the right time
2) Validate the developed approaches and algorithms using independent data from an operating plant.

6. ACKNOWLEDGMENTS

This report was made possible through funding from the U.S. Department of Energy's Office of Nuclear Energy under the Nuclear Energy Enabling Technologies Program. We are grateful to Melissa Bates at the U.S. Department of Energy and Pattrick Calderoni at Idaho National Laboratory for championing this effort. The data used in this research was provided by Exelon Generation Company. We are also thankful to the plant engineers for their technical discussions on the data.

7. REFERENCES

[1] R. Bryce. 2020. "Nuclear Plants In Illinois Are Slated For Closure. Will the State's Democratic Politicians Save Them?" Forbes, Aug 31, 2020. https://www.forbes.com/sites/robertbryce/2020/08/31/nuclear-plants-in-illinois-are-slated-for-closure-will-the-states-democratic-politicians-save-them/?sh=5e6cf62511ef.

[2] S. M. Bragg-Sitton, R. Boardman, C. Rabiti, and J. O'Brien. 2020. "Reimagining Future Energy Systems: Overview of the US Program to Maximize Energy Utilization via Integrated Nuclear-Renewable Energy Systems." International Journal of Energy Research 44, no. 10 (August 2020): 8156–8169. First published 2020 Feb 17. https://doi.org/10.1002/er.5207.

[3] Y. Lei, B. Yang, X. Jiang, F. Jia, N. Li, and A. K. Nandi. 2020. "Applications of Machine Learning to Machine Fault Diagnosis: A Review and Roadmap." Mech. Syst. Signal Process. 138: 106587. https://doi.org/10.1016/j.ymssp.2019.106587.

[4] Nuclear Energy Institute. 2011. "Industry Guideline for Monitoring the Effectiveness of Maintenance at Nuclear Power Plants." NUMARC 93-01, Revision 4A. nrc.gov/docs/ML1111/ML11116A198.pdf.

[5] K. A. Manjunatha and V. Agarwal. 2019. "Techno-Economic Analysis Framework for Wireless Networks in Nuclear Power Plants." INL/EXT-19-55830, Rev 0, Idaho National Laboratory.

[6] K. A. Manjunatha and V. Agarwal. 2020. "ISM Band Integrated Distributed Antenna Systems for Industry 4.0: A Techno-Economic Analysis." GLOBECOM 2020 - 2020 IEEE Global Communications Conference, Taipei, Taiwan, pp. 1–6, https://doi.org/10.1109/GLOBECOM42002.2020.9322617.

[7] V. Agarwal. 2020. "Wireless Sensor Modalities at A Nuclear Power Site to Collect Vibration Data." INL/EXT-20-58548, Rev 0, Idaho National Laboratory.

[8] Pentasense. 2019. "Vibration Mote." Product page. Accessed May 27, 2020. https://petasense.com/products/vibration-motes/.

[9] Bently Nevada. 2020. "Bently Nevada Ranger Pro Wireless Vibration Sensor." Product page. Accessed May 27, 2020. https://www.bakerhughesds.com/bently-nevada/online-condition-monitoring/ranger-pro-wireless-sensor-system

[10] KCF Technologies. 2017. "Vibration Sensor Node (SD-VSN-3)." Product page. Accessed May 27, 2020. https://www.kcftech.com/smartdiagnostics/products/sensors.html

[11] KCF Technologies. 2015. "SmartDiagnostics® Application Note: Wireless Interference and SmartDiagnostics®." PDF. Published May 27, 2015. https://www.kcftech.com/smartdiagnostics/resources/whitepapers/Application%20Note%20-%20Wireless%20Interference%20and%20SmartDiagnostics.pdf

[12] Bently Nevada. 2020. “Bently Nevada Proximity Sensors.” Product page. Accessed May 27, 2020. https://www.bakerhughesds.com/bently-nevada/sensors/proximity-sensors.

[13] C. Walker, P. Ramuhalli, V. Agarwal, and N. Lybeck. 2021. “Nuclear Power Fault Diagnostics and Preventative Maintenance Optimization.” INL/EXT-21-61427, Rev 0, Idaho National Laboratory.

[14] P. Ramuhalli, C. Walker, V. Agarwal, and N. Lybeck. 2020. “Development of Prognostic Models using Plant Asset Data.” ORNL/TM-2020/1697, Oak Ridge National Laboratory.

[15] A. Oluwasegun and J. C. Jung, “The Application of Machine Learning for the Prognostics and Health Management of Control Element Drive System.” Nucl. Eng. Technol. 52, no. 10: 2262–2273. https://doi.org/10.1016/j.net.2020.03.028.