

Dashboard: Nonintrusive Electromechanical Fault Detection and Diagnostics

Daisy Green, Peter Lindahl
& Steven Leeb
Massachusetts Institute of Technology
Cambridge, MA 02139
Email: dhgreen@mit.edu

Lt. Thomas Kane & Lt. Stephen Kidwell
U.S. Coast Guard
& Massachusetts Institute of Technology
Cambridge, MA 02139

John Donnal
U.S. Naval Academy
Annapolis, MD 21402

Abstract—Modern power monitoring systems record vast amounts of equipment operational data. For these systems to improve efficiency and performance, the data must be presented as an intuitive decision aid for watchstanders. The Nonintrusive Load Monitor (NILM) dashboard provides actionable information for energy scorekeeping, activity tracking, and equipment fault detection and diagnostics (FDD). Electrical monitoring through the NILM dashboard can identify both “soft” faults (the gradual degradation of equipment performance) and “hard” faults (the complete failure of a piece of equipment). This paper presents metrics and visualizations that have proven useful for FDD. Analysis is presented from case studies of the NILM dashboard for identifying fault conditions aboard two United States Coast Guard cutters (USCGCs), SPENCER and ESCANABA.

Index Terms—Condition-based maintenance, energy efficiency, fault detection, Nonintrusive load monitoring

I. INTRODUCTION

There are at least three philosophies for maintaining mission-critical systems. The most basic form of maintenance is Corrective Maintenance (CM), where equipment is run until failure and then replaced. Failures or “casualties” are corrected when identified, but may impact the operational availability of the system. To maximize operational availability, many organizations elect to use a Preventative Maintenance (PM) schedule. PM seeks to maximize equipment service life by conducting maintenance actions at set periodic intervals, regardless of the actual health of the system. The drawback to the PM approach is that time-consuming maintenance activities are often repeatedly performed on healthy equipment. Furthermore, even with a PM program in place, sudden equipment failures can still occur without warning. A third option, condition-based maintenance (CBM), seeks to resolve the inefficiencies of CM and PM by conducting maintenance activities based on data gathered from condition-monitoring of equipment [1]. Preventative steps are only taken when there is evidence of degraded equipment health. Data is used to provide advance warning of failure so that equipment replacement can be scheduled [1]. There is an even greater incentive to implement CBM for deployable units, where the impact of a single equipment failure can severely limit mission effectiveness. The logistics of acquiring a replacement part while in the field further multiplies costs and asset downtime. An effective CBM program maximizes the maintenance conducted during

scheduled availabilities and minimizes sudden failures when deployed.

There are numerous advantages to condition-based maintenance programs, but there are also several significant hurdles to implementation. First, healthy equipment behavior can be subjective. Due to variations in environment and application, data gathered from lab testing or from manufacturer’s specifications often does not match the parameters measured on healthy equipment in the field. An accurate baseline can be established by continually monitoring thousands of systems across a broad variety of use cases but many organizations that can benefit from a CBM program will not have access to such a complete operational data set. Furthermore, true indicators of impending equipment failure are unknown until a failure occurs while monitoring is in place. Establishing warnings for a failure mode that has not yet been observed requires complex causal-model algorithms, that often need significant computing resources and are unlikely to be trusted by the end user without understanding the underlying physics [2]. Reliable failure data can also be obtained by setting up a test facility to run equipment to failure in a variety of modes, but this is also an unreasonable undertaking for most organizations.

Power monitoring can still be used for a condition-based maintenance program, even without a known baseline or complete set of failure data. Electrical measurements can be used to monitor equipment behavior and deviations from acceptable or standard condition. These deviations may be the result of a “soft fault”, where equipment continues to run, but with increased wear and energy consumption that if left undetected eventually leads to complete failure. A soft fault warning can be given when a parameter reaches a threshold deviation from the initial condition. The warning does not mandate that equipment must be repaired or replaced, but serves to draw attention to potential failures. These warnings are highly valuable because soft faults are often invisible without condition-monitoring in place. The challenges of soft fault detection, the resulting increase of energy consumption and component wear, and the decrease in system efficiency are highlighted in Reference [3]. The vast implementation of automatic controllers and feedback control often masks soft faults by automatically altering run-times and drawing excess power to maintain system operability. By choosing the appropriate parameters and setting reasonable warning




Fig. 1. Main engineering deck of USCGC SPENCER.

thresholds, a fault detection and diagnostics (FDD) approach to condition-based maintenance can be implemented without a large data set or up-front investment.

Fundamental to any fault detection and diagnostics system is the data and the understanding of the data. Many power monitoring systems record vast amounts of equipment operational data, but users are not able to understand the relationship between the data and equipment failures [4]. For a successful FDD program, the data must be presented as an intuitive decision aid for users. This paper presents metrics and visualizations that have proven useful for a FDD based CBM program implemented with Nonintrusive Load Monitoring aboard two United States Coast Guard cutters (USCGCs), SPENCER and ESCANABA.

II. NONINTRUSIVE LOAD MONITORING ONBOARD US COAST GUARD CUTTERS

The USCGCs SPENCER and ESCANABA are Famous class, 270 ft (82 m) medium endurance cutters (MECs), based in Boston, MA. The ships each maintain a 100 person crew and an operational tempo of 180 days at sea per year, performing a host of Coast Guard missions, including environmental stewardship, law enforcement, fisheries protection, and national security.

On legacy ships, such as these MECs, watchstanders manually record readings from equipment at local gauges and panels throughout the ship. Machinery control and monitoring systems (MCMS) are either installed in limited areas or not presented at all. Even though these legacy cutters lack a fully integrated MCMS, they still contain many closed-loop, automated systems. Closed-loop systems are actuated by sensor feedback, such as tank-level indicators or temperature sensors. These sensor feedback systems are crucial to equipment operations, but they make faults difficult for watchstanders to detect. For instance, if there is a vacuum leak in the sewage system, the system will still draw vacuum and the toilets will continue to flush, but the pump will run much more often. Faults of this nature are nearly invisible to watchstanders, but are clearly visible in the electric power readings.



Fig. 2. NILM installation on SPENCER, containing a data acquisition unit and host computer in a single box. Cables connect to the the electrical panel where current and voltage sensors are located.

Nonintrusive Load Monitor (NILM) systems are installed upstream of two 440V sub-panels in the main engine room, on both SPENCER and ESCANABA. These panels, STBD (3-117-1) and PORT (3-117-2), supply a variety of mission-critical systems, crucial to the proper operation of ship propulsion, power generation, and auxiliary services. These panels supply the loads that support the twin ALCO V18 Main Propulsion Diesel Engines (MPDE) propelling the ship, as well as the two 475 kW V12 Caterpillar ship-service diesel generators (SSDG) providing power to the ship's microgrid. Figure 1 shows the main engineering deck of the USCGC SPENCER, with the two MPDE located on either side of the deck plates.

Nonintrusive Load Monitor (NILM) systems are installed on these vessels to serve two primary objectives: first, to identify equipment operating schedules to improve watchstander situational awareness; then, once the operating schedules have been accurately identified, analyze the gathered data for condition-based maintenance and fault detection to improve system operational availability. A NILM meter consists of a data acquisition unit (DAQ) that samples current and voltage sensor outputs at 8 kHz and transmits the data via Ethernet to a host computer. Figure 2 shows a NILM installation that contains a DAQ and host computer in a single box. The computer converts the current and voltage data into 60 Hz real power (P), reactive power (Q), and higher harmonic content using the Sinefit algorithm [5]. P and Q correspond to the envelopes of in-phase and quadrature current drawn relative to the load voltage [6]. The power information directly corresponds to the physics that governs load behavior, producing distinct signatures for each piece of equipment. The NILM is trained to recognize the signatures from previous shipboard observation. This allows the NILM to detect load *events*, which occur when equipment transitions between *on* and *off* states, thus generating the operating schedules of individual equipment. Then, a NILM system can determine load *metrics*, or diagnostic indicators, which are statistical

conclusions that expose anomalies and patterns to be used for CBM. However, the value of NILM as a CBM tool depends on the appropriate choice of diagnostic indicators and fault warning levels. Sections III and IV describe the selected parameters for this case-study and the process of determining fault warning levels, respectively. Then, Section V shows the utility of these techniques by using the NILM dashboard graphical user interface to present *metrics* to crewmembers and identify fault scenarios aboard SPENCER and ESCANABA.

III. DIAGNOSTIC INDICATORS FOR FAULT DETECTION

Symptoms of impending equipment failure are often present in the electrical data long before a breakdown occurs. Non-intrusive load monitoring records both the electrical signature and the operating schedule for a piece of equipment, allowing for a broad range of fault diagnostic methods. It is crucial to select the appropriate parameters for condition-monitoring to create a useful tool that provides actionable information for watchstanders. For this study, the following five parameters were selected for equipment diagnostics:

- Power: Steady-state real power.
- Power Factor: The ratio of real power to apparent power.
- Average Run Duration: Time between activation and shutdown.
- Total Run Time: Total time the equipment is online over a 24-hour period.
- Daily Actuations: Number of discrete operations per day.

These parameters work well for the equipment monitored onboard the MECs, but other metrics may be useful in other environments. Regardless of the criterion selected, an effective CBM program using NILM should track changes in both equipment signature and equipment behavior. For this study, *power*, *power factor* and *average run duration* track the equipment signature. These metrics can detect material degradation of equipment, such as mechanical wear and corrosion. The *average run duration*, *total run time*, and *daily actuations* track equipment behavior, and are used to find sensor and automation faults that might cause equipment to run too frequently or not enough.

The selected parameters were tailored for the equipment served by the monitored sub-panels. For example, the equipment in this work consists largely of pumps and heaters, each with a consistent steady-state power signature. Therefore, tracking the steady-state *power* and *power factor* can detect equipment wear or changes in operational behavior. A change in *power* demand may indicate a worn motor bearing [7] or a change in *power factor* could be a sign of corroded heated elements. In other situations, where variable speed drives or other non-linear loads are present in the system, the harmonic content of the electrical signatures may be an appropriate diagnostic indicator [8]. Many of the heaters and pumps monitored by NILM on the MECs are controlled by closed-loop automated systems such as tank-level sensors or thermostats. Therefore, it is useful to track the *average run duration*, *daily actuations* and *total run time* to detect broken sensing systems. A fouled tank level indicator or failed




Fig. 3. Normal sequence of operations when a MPDE engine secures. Pre-lube pump activates and then the lube oil heater activates shortly after. The presence of one load without the other can be used as a fault warning.

thermostatic controller can, for example, cause equipment to activate in repeated short-cycles or run for excessively long periods [9].

Conversely, a lack of equipment operation can often indicate a sensor fault or even a complete failure of a pump or heater. In the event a load is no longer operating, or if a load's signature has changed to the point where accurate load identification is impossible, the *daily actuations* parameter can still alert the user that a fault is present. For instance, the graywater pump system should run at minimum once every 3-4 hours, based on healthy system data. If the NILM does not detect a graywater pump activation for an extended period of time, this could be the indication of a fault. The 24-hour period to observe run frequency can be tailored to the equipment being monitored to more rapidly detect a fault detection, i.e., *hourly actuations*.

The interdependence of equipment can also be used to diagnose fault conditions, and was useful in locating the MPDE jacket water heater fault detailed later in Section V. Finite State Machine (FSM) loads have multiple stages of operation that must occur in the same order during normal operations [10]. For instance, the MPDE lube oil heater should not be run unless the pre-lube pump is already online. Figure 3 shows the proper order of operations when the MPDE is secured. The presence of one load online without the other can be evidence of a fault condition.

Finally, it is important to note that a single extended pump run or even a few frequent runs is not necessarily a cause for concern. These may occur during manual operation or maintenance. This is accounted for by tracking the *average* over 24 hours for averaged parameters, i.e. *power*, *power factor*, and *average run duration*, and the *total* over 24 hours for summated parameters, i.e. *total run time* and *daily actuations*. The 24-hour window serves to help prevent the gauges from falsely displaying an alarm as the result of a brief anomaly. The 24-hour period can easily be adjusted for different applications where loads activate less or more frequently or tighter controls are required.

IV. DETERMINING FAULT WARNING LEVELS

Condition-based maintenance parameters are communicated on the NILM dashboard via “green-yellow-red” diagnostic gauges. The green region represents healthy operation, while the yellow region can be considered analogous to a trouble warning, and the red region is a more definitive fault alarm.

Determining the proper threshold for each region on the gauges is crucial to making a useful, actionable tool for the ship's crew. A variety of methods have been proposed to determine fault thresholds for industrial applications [2]. For this study, a statistical process control (SPC) method is implemented. Effective SPC attempts to differentiate between natural variations and variations that are due to process failure [11].

Historical data collected by NILM was used for SPC analysis. *Metrics* were calculated for NILM data collected from the SPENCER and ESCANABA sub-panels dating back to 2016. Equipment nameplate information and the ship's logs were used to exclude any data that may have been from a previous fault condition. Deviation from the historical data for any parameter is evidence of a possible fault. SPC provides a method to determine exactly how much deviation is acceptable and when a deviation should trigger a fault warning. The SPC method consists of determining a centerline, an upper control limit (UCL) and a lower control limit (LCL). Warnings are issued when a parameter reaches the upper or lower control limits.

First consider a continuous variable, in which the variable can fall anywhere within a particular range of values, such as *power*, *power factor*, *average run duration*, and *total run time*. Considering the standard normal distribution, SPC uses the arithmetic mean (θ) of the parameter as the centerline [11]. The UCL and LCL are defined as,

$$UCL = \theta + k\sigma \quad (1)$$

$$LCL = \theta - k\sigma \quad (2)$$

where σ is the standard deviation and k is an integer that sets the distance of the control limits. For a parameter with a normal distribution, $k = 3$ is the accepted industry standard for a fault warning [11], [12], and corresponds to the red region on the gauges. The choice of control limits affects the risks of Type I or Type II errors, where Type I errors refer to incorrectly reported faults and Type II errors refer to missed faults. By widening the control limits, the risk of Type I errors decreases; however there is an increased risk of Type II errors as more data points will fall within the control limits and viewed as normal. Contrarily, if the control limits are narrowed, there is an increased risk of Type I errors and decreased risk of Type II errors, as more data points will fall outside the control limits and classified as fault conditions. The "3-sigma" rule is conservative and designed to minimize the risk from false alarms. However, analysts often suggest using two sets of limits; *action limits* at "3-sigma" and *warning limits* at "2-sigma" [11]. For this application the intermediate control at $k = 2$ corresponds to the yellow region on the gauge. Addition of the intermediate control limit provides more rapid detection of faults [13]. Figure 4 shows how SPC maps the probability density function (PDF) of a normal distribution to the green, yellow and red regions of the gauges. The percentages in each region correspond to the




Fig. 4. Probability density function of a normal distribution showing progressive thresholds for fault detections. Colors correspond to the red, yellow, and green regions on the dashboard gauges.

likelihood that the variable falls within that particular range of values. The quantile values displayed at the bottom of Figure 4 correspond to the probability that some variable, X , is less than or equal to some value, x , where x is the centerline and control limits. This can be written as a cumulative distribution function (CDF),

$$F(x) = Pr(X \leq x). \quad (3)$$

The inverse cumulative distribution function (ICDF), or quantile function,

$$F^{-1}(p) = x, \quad (4)$$

solves for the x value that would make $F(x)$ return some probability, p .

The SPC process can be adapted if the normal distribution does not properly fit the data. For example, the Weibull distribution is often used in machinery reliability applications [12]. The PDF for a two-parameter Weibull function is:

$$f(x) = \left(\frac{\beta}{\alpha}\right) \left(\frac{x}{\alpha}\right)^{\beta-1} e^{-\left(\frac{x}{\alpha}\right)^\beta} \quad (5)$$

where α is the scale parameter and β is the shape parameter. To create the gauge regions for a non-normal distribution, the probability quantiles (p) should match the red, yellow, and green regions of the normal distribution in Fig. 4.

The ICDF function for a Weibull distribution is:

$$x = F^{-1}(p|\alpha, \beta) = \alpha(-\ln(1-p))^{1/\beta} \quad (6)$$

Therefore, the centerline can be found by setting p equal to .50 and solving for x . The upper and lower yellow threshold levels can be found by setting p to 0.977 and 0.023, respectively. Similarly, the upper and lower red threshold levels can be found by setting p to 0.999 and 0.001, respectively. This ensures that the probability of an alarm detection is the same regardless of the PDF selected for modeling. Each parameter monitored by the NILM dashboard can be analyzed individually and the gauges adjusted to provide diagnostic warnings at appropriate levels.

Next consider a discrete variable, in which the variable has finite values, such as the number of *daily actuations*. Because these can only occur as integer values, the CDF is not continuous and increasing; thus the generalized inverse




Fig. 5. Different probability density functions fit to each parameter of the controllable pitch propeller (CPP) pumps.

distribution function will be used instead of the ICDF. The generalized inverse distribution function can be expressed as,

$$x = F^{-1}(p) = \inf\{x \in \mathbb{R} : F(x) \geq p\}, \quad (7)$$

where \inf is the infimum, or the greatest lower bound. Similar to a continuous function, centerline and upper and lower limits can be determined by setting p to the appropriate values and solving for x .

The statistical process control method is visualized in Figure 5 for the MEC CPP pumps. Histograms for each metric are created using historical NILM data. Then, probability densities are fit to the data. The probability densities visualized in Figure 5 are normalized to match the total area of the histograms. For continuous functions, the histograms are modeled with multiple probability density functions (PDF), while a discrete function is modeled with probability mass functions (PMF), and the best fit function is selected. For continuous functions, the Anderson-Darling (AD) test returns a decision for the null hypotheses that the data is from a population with a specific distribution [14]. The test rejects the null hypothesis at the 5% significance level. For the example shown in Figure 5, the AD-test did not reject the null hypothesis for a normal distribution for *power* and *power factor*, with *p-values*, or probability values, of 0.2446 and 0.6561, respectively. The AD-test did not reject the null hypothesis for a Weibull distribution for *average run duration* and *total run time* with *p-values* of 0.4430 and 0.7205, respectively. To test the discrete models, the chi-square goodness-of-fit test is used since it is applicable for discrete distributions [14]. The chi-square test did not reject the null hypothesis for a Poisson distribution for number of *daily actuations* with a *p-value* of 0.0511. Note that the model for *daily actuations* is only for *if* the load ran that day. If there were no runs for a given day, that is demonstrated with a grayed out gauge. Depending on the PDF or PMF selected, the fault detection thresholds are set as shown in Figure 4 or using a quantile function as described in Equations (6) and (7). This process was repeated to determine control limits for each of the monitored loads.

V. NILM DASHBOARD

The on-site NILM dashboard is an easy-to-understand visual display that provides users with useful insight into the




Fig. 6. The NILM dashboard display on USCGC SPENCER.

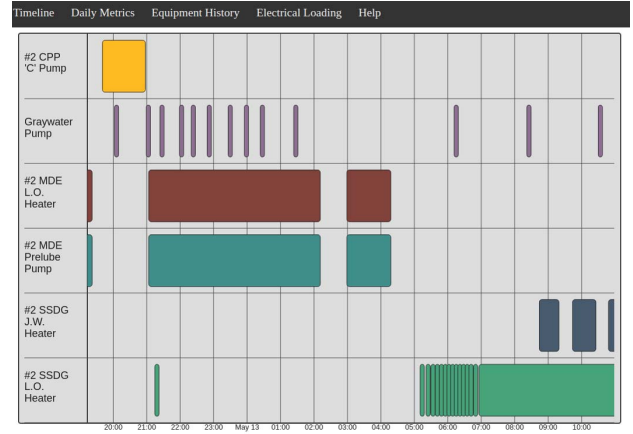


Fig. 7. The Timeline view of the NILM dashboard displays equipment status over a day of at sea operation. Colored blocks indicate periods where equipment is energized.

operation of the electrical system and its loads. The interface uses the diagnostic metrics on system health, as discussed in Section III, to present both real-time and historic equipment information to watchstanders [15]. On SPENCER, the NILM dashboard is conveniently installed in the engine room, servicing both PORT and STBD loads, as shown in Figure 6. The information is made available to crew through multiple interactive tools, each serving a unique objective. First, the dashboard processes information about load *events*, through the “Timeline” view. The Timeline displays the live equipment status and recent operation of all monitored loads. This tool aids in improving situational awareness by allowing a crew member to quickly deduce the recent operating schedule of any of the monitored loads. An example of the Timeline from a day of underway operation is shown in Figure 7. Next, the NILM dashboard maintains load *metrics*, implemented through *rolling* metrics over the past 24-hour window and *daily* load metrics for each day, ending at midnight. These metrics, which are described in Section III, are presented to the crew members through the “Metrics” and “Historic” views. The Metrics and Historic views serve as the primary tools for condition-based maintenance. The Metrics view displays the *rolling* metrics with “green-yellow-red” gauges, using the fault warning levels described in Section IV. The Historic view supplements the analysis of the Metric view by presenting a bar graph dis-

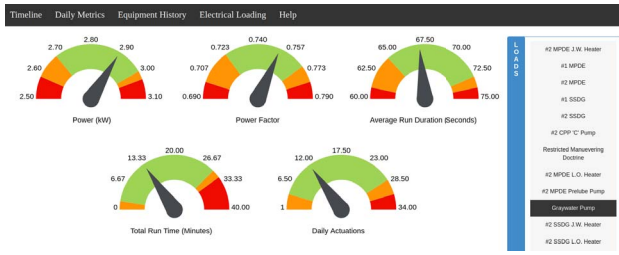


Fig. 8. NILM dashboard Metrics view indicating healthy graywater pump operation.




Fig. 9. NILM dashboard Historic view showing healthy graywater pump run frequency (*daily actuations* metric).

playing *daily* metrics for a user selected piece of equipment and diagnostic metric, tracking equipment behavior over time. The Metrics and Historic views for healthy graywater pump operation are shown in Figure 8 and Figure 9 respectively. Together, these interfaces can serve to direct watchstanders to previously invisible faults within the system.

A. MPDE Jacket Water Heater

The MPDE jacket water (JW) heater consists of two 4.5 kW heating elements on either side of the engine block, shown in Figure 10. The heaters appear to a NILM as a single electrical load, as a 9kW resistive load with 3kW per phase. The controller configuration results in the MPDE JW heater frequently turning *on* or *off* simultaneously with the MPDE lube oil (LO) heater and both the LO heater and MPDE pre-lube (PL) pump. The PL pump and LO heater work in tandem and are served from the same controller. The controller is activated automatically when the engine speed falls below 150 RPM. When the controller is activated, the PL pump comes online and the LO heater enters automatic mode. In automatic mode, the heater energizes when the oil temperature falls below 90 F and secures when the oil temperature goes above 120 F. If the controller is activated when the engine is already cold, the PL pump and LO heater will activate simultaneously. If the engine is warm, the PL pump will run by itself until the engine temperature falls below the threshold and the LO heater comes online. The JW heater is activated from a separate controller and thermostat. However, the JW heater also runs automatically based on the 90 F and 120 F setpoints, and lube oil and jacket water temperature normally track closely together. Figure 11 shows the electrical transients for a healthy MPDE JW heater from SPENCER operating alone and in tandem with the other MPDE system loads.




Fig. 10. Left: Jacket water heater mounts to the engine block between the cylinders. Right: New MPDE jacket water heater.




Fig. 11. *On* transient of MPDE (a) pre-lube pump, lube oil heater, and jacket water heater, (b) lube oil heater and jacket water heater, and (c) jacket water heater.

On the NILM dashboard, healthy heater operation is captured on the Metrics gauges, shown in Figure 12. As expected, normal operation is a 9kW load with a nearly zero reactive power. Heater behavior is characterized by runs of 15 minutes to 4 hours, no more than a few times per day. However, as seen in Figure 13, the NILM detected a slow change in the heater's electrical signature. Over the course of more than a year, the heater's steady-state real power draw decreased in a series of steps starting at 3000 W per phase (normal operation) until it eventually reached 1500 W per phase. During the same period, the heater, normally a purely resistive load, showed an increase in reactive power on different phases of power.

It is important to consider how a degrading electrical signature affects the load classification process. For the NILM dashboard to be an effective fault diagnostic tool, changing electrical transients such as Figure 13 must be properly identified and tracked. The artificial neural network (ANN) approach used to identify loads on the SPENCER and ESCANABA relies on supervised learning where load transients are hand-labeled in order to train the identification algorithm [16], [17]. These hand-labeled transients are healthy loads and the ANN may incorrectly classify damaged equipment as another load unless trained on changes that might occur during a fault. However, despite the challenges presented by the changing load signature, the NILM dashboard was still able to effectively identify the faulty MPDE JW heater with no adjustment to the identification algorithm. The dashboard is effective in diagnosing this fault because of the simultaneous turn-on with the MPDE LO heater and PL pump as previously described.




Fig. 12. NILM dashboard showing a healthy MPDE jacket water heater condition. All gauges are within the green regions.

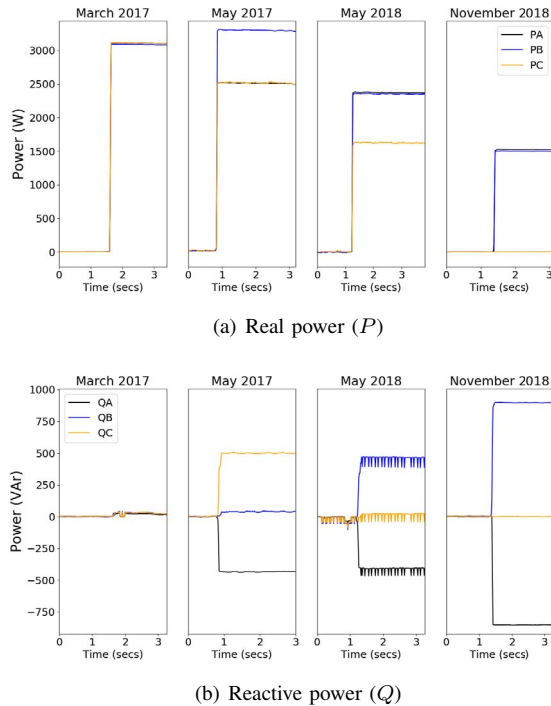


Fig. 13. On transients of (a) real power and (b) reactive power for MPDE jacket water heaters, observed over more than a year on the port engine on USCGC SPENCER.

Likewise for turn-off transients, the MPDE JW heater and LO heater often turn-off simultaneously. The combined actuations reduced the effect of the diminished power demand on the classification algorithm because there are no other loads nearby in the feature space. Figure 14 shows the Metrics gauge tracking *power* moving from the green region into the yellow at the first sign of abnormality (May 2017). This indicator should prompt a watchstander to check and inspect the heaters. Once the heater reaches the state shown in May 2018 in Figure 13, the real power draw has decreased enough that the gauge enters the red region, indicating that a failure has likely occurred.

The nature of this fault can be applied to improve future fault detection. For example, the ANN can be trained to identify the degraded transients in Figure 13 as JW heater activations. The algorithm can also be trained to look for similar decay patterns in other heaters in the system, i.e. a small reduction in real power demand coupled with the appearance of a reactive component. This fault also demonstrates




Fig. 14. The *power* gauge from the NILM dashboard captures the degradation of a MPDE jacket water heater as it transitions from a (a) healthy system, then a (b) potentially degraded system and finally a (c) degraded system.



Fig. 15. Left: Bare wiring found in SPENCER port inboard jacket water heater. Right: Jacket water heaters removed from engines after electrical degradation was observed, revealing severe corrosion (top from ESCANABA port inboard heater, removed April 2019, and bottom from SPENCER port inboard heater, removed December 2018).

the importance of having a parameter that can detect the lack of operation as a fault. If for example, the MPDE JW did not turn on or off simultaneously with the LO heater, it may no longer be recognizable. In that scenario, the *daily activations* metric would drop to zero. If a crewmember knows that the load *should* be running, that will alert them to a possible fault. Using the interdependency of the MPDE system as an example, if a LO heater activation is repeatedly detected without a JW heater activation, this would be a flag for a fault condition.

The electrical signature degradation in Figure 13 was observed on both SPENCER and ESCANABA. After alerting the ships' crews, subsequent removal and inspection of the heaters revealed severe corrosion of the heating elements. Figure 15 shows the physical condition of the heaters after their electrical signature decayed. The damage to the heating elements was also allowing jacket water to enter the heater's electrical enclosure. Upon removal of the enclosure covers, the wiring was degraded and one of the heaters was lightly smoking (see Figure 15). In addition to identifying a failed heating element, NILM analysis detected a potential shock hazard and likely prevented a Class Charlie electrical fire in the engine room.

For a variety of reasons, this fault was nearly impossible to detect without the assistance of NILM. Despite the holes in the heaters shown in Figure 15, there was no ground detected. The stray current flowed into the jacket water and did not reach the ship's hull. The heater controller showed that the heaters were online and the thermostatic controller continued




Fig. 16. Conceptual diagram of the graywater system [9].

to activate the load. There is currently no PM action that prompts the crew to check the heater enclosures for damage or circuit continuity. Like most soft faults, detecting this issue with normal watchstanding efforts is an unrealistic expectation. The root cause of the corrosion is unknown at this time, but it is clear that this is a systemic issue. Of the eight heaters in place on ESCANABA and SPENCER, seven were observed to have suffered complete failure or were operating with reduced heating capability, in less time than their rated service life. A memorandum regarding this failure was sent to US Coast Guard MECPL (Medium Endurance Cutter Product Line) for increased visibility.

B. Graywater Pumps

The graywater disposal system collects, transfers and disposes of the relatively clean water from showers, sinks, washing machines and other appliances. The water is stored in a 138 gallon (522 L) holding tank and pumped overboard or to a larger storage tank when full. Two identical pumps (for redundancy) alternate each cycle to discharge the tank. Conductivity sensors detect water levels and provide feedback control to the pumps. A pump turns on and starts discharging when water reaches the high level sensor (92 gallons) and turns off when water reaches the low level sensor (13 gallons). Thus, the pump expels 79 gallons (300 L) during each normal pump run. A conceptual diagram of the graywater system is shown in Figure 16.

During monitoring of SPENCER, the gray water system experienced two types of faults, the first a failed high-level tank sensor and the second a failed check valve [9]. If either of these fault conditions goes undetected, it could easily result in pump failure and the loss of a crucial auxiliary system. These faults demonstrate the ability of the *average run duration*, *total run time*, and *daily actuations metrics* to diagnose fault conditions.

1) *Faulty High-Level Sensor*: The high-level sensor failed in December 2015 due to solid residue shorting its electrodes, resulting in constant high-level readings. This caused the pumps to short-cycle. This fault went unnoticed by the crew but could be easily detected with the NILM dashboard. The




Fig. 17. Metrics interface displaying faulty high-level sensor fault condition of the graywater pumps.




Fig. 18. Metrics interface displaying failed check valve fault condition of the graywater pumps.

graywater events were so short, typically only a few seconds, that the load identifier was only able to identify a fraction of the number of actual graywater events. Many events were less than a second because the tank was essentially empty, which left the transient unidentifiable. Even so, the number of *detected* events on some days is up to 10 times greater than normal. Figure 17 shows the Metrics view for a day displaying this fault condition. As expected, the *average run duration* is much shorter than normal and the number of *daily actuations* is much greater than normal. As a result, these metrics are in the red region.

2) *Failed Check Valve*: From January to March 2017 the graywater system exhibited another fault condition due to a broken check valve. During normal operation, graywater drains from throughout the ship to the holding tank and the monitored pumps transfer this to a larger storage tank. However, the broken check valve allowed water to flow backwards from the storage tank to the holding tank. The pumps had to run almost continuously to keep the holding tank from overflowing. This fault condition can also be detected with the NILM dashboard Metrics view as shown in Figure 18. As expected, the *daily actuations* and *total run time metrics* are much greater than normal operation, and are deep in the red region. Due to the almost continuous flow of water into the holding tank, the pump would have to run for longer as it would attempt to empty the tank. As a result, the *average run duration* metric is also in the red region.

VI. CONCLUSION

These case-studies demonstrate the effectiveness of NILM and the importance of a condition-based maintenance program. Identifying faults and degraded capability prior to the

ship's deployment can increase the crew's effectiveness and maximize operational availability. NILM represents a low-cost method for the rapid implementation of a condition-based maintenance and fault detection program for electromechanical equipment. The failures in the MPDE JW heater and graywater pumps show that through careful selection of metrics and simple statistical analysis, NILM can quickly detect a broad range of system anomalies and assess individual equipment health. The NILM dashboard brings this crucial information to the field, ensuring that operators have real-time, actionable information on mission-critical systems.

ACKNOWLEDGMENT

This work was funded by The Grainger Foundation and the Office of Naval Research NEPTUNE program. We gratefully acknowledge the support and dedication of the US Coast Guard, and, in particular, the spectacular crews of USCGC ESCANABA and SPENCER. Additional support for this work was provided by the Cooperative Agreement between the Masdar Institute of Science and Technology (Masdar Institute), Abu Dhabi, UAE and the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA - Reference 02/MI/MIT/CP/11/07633/GEN/G/00.

REFERENCES

- [1] A. Jardine, D. Lin, and D. Banjevic, "A review on machinery diagnostics and prognostics implementing condition-based maintenance," *Mechanical Systems and Signal Processing*, vol. 20, no. 7, pp. 1483–1510, Oct 2006.
- [2] V. V. Sourabh Dash, "Challenges in the industrial applications of fault diagnostic systems," *Computers and Chemical Engineering*, vol. 24, pp. 785–791, 2000.
- [3] M. Breuker and J. Braun, "Common faults and their impacts for rooftop air conditioners," *HVAC&R Research*, vol. 4, no. 3, pp. 303 – 318, Jul 1998.
- [4] X. Dai and Z. Gao, "From model, signal to knowledge: A data-driven perspective of fault detection and diagnosis," *IEEE Transactions on Industrial Informatics*, vol. 9, no. 4, pp. 2226 – 2238, Nov 2013.
- [5] J. Paris, J. S. Donnal, Z. Remscrim, S. B. Leeb, and S. R. Shaw, "The sinefit spectral envelope preprocessor," *IEEE Sensors Journal*, vol. 14, no. 12, pp. 4385–4394, 2014.
- [6] C. Laughman, K. Lee, R. Cox, S. Shaw, S. Leeb, L. Norford, and P. Armstrong, "Power signature analysis," *IEEE Power and Energy Magazine*, vol. 1, no. 2, pp. 56–63, Mar 2003.
- [7] R. Beebe, "Estimate the increased power consumption caused by pump wear," *Pump Magazine*, vol. 58, pp. 20–27, 2008.
- [8] K. D. Lee, S. B. Leeb, L. K. Norford, P. R. Armstrong, J. Holloway, and S. R. Shaw, "Estimation of variable-speed-drive power consumption from harmonic content," *IEEE Transactions on Energy Conversion*, vol. 20, no. 3, pp. 566–574, Sept 2005.
- [9] P. A. Lindahl, D. H. Green, G. Bredariol, A. Aboulhian, J. S. Donnal, and S. B. Leeb, "Shipboard fault detection through nonintrusive load monitoring: A case study," *IEEE Sensors Journal*, vol. 18, no. 21, pp. 8986–8995, Nov 2018.
- [10] G. W. Hart, "Nonintrusive appliance load monitoring," *Proceedings of the IEEE*, vol. 80, no. 12, pp. 1870–1891, 1992.
- [11] D. C. Montgomery, *Introduction to statistical quality control*. John Wiley & Sons, Inc., 2009.
- [12] S. Rao, *Reliability-Based Design*. McGraw-Hill, 1992.
- [13] W. A. Levinson, *Statistical process control for real-world applications*. CRC Press, 2011.
- [14] M. J. de Smith, *Statistical Analysis Handbook*. The Winchelsea Press, Drumlin Security Ltd., 2018.
- [15] A. Aboulhian, D. Green, J. Switzer, T. Kane, G. Bredariol, J. Donnal, P. Lindahl, and S. Leeb, "Nilm dashboard: A power system monitor for electromechanical equipment diagnostics," *IEEE Transactions on Industrial Informatics*, vol. 15, no. 3, pp. 1405–1414, Jun 2018.
- [16] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 05 2015.
- [17] D. Green, S. Shaw, P. Lindahl, T. Kane, J. Donnal, and S. Leeb, "A multi-scale framework for nonintrusive load identification," *IEEE Transactions on Industrial Informatics*, Jun 2019.