NILM Dashboard: Actionable Feedback for Condition-Based Maintenance

Daisy Green, Thomas Kane, Stephen Kidwell, Peter Lindahl, John Donnal, and Steven Leeb

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 condition-based maintenance (CBM). Using a NILM to present metrics that track changes in *equipment signature* and *equipment behavior* allows for effective CBM. 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 CBM. Analysis from case studies of fault conditions identified aboard two United States Coast Guard cutters (USCGCs), SPENCER and ESCANABA, are discussed.

Goals of Condition-Based Maintenance

As the complexity and criticality of power systems increase, so does the need for condition- based maintenance (CBM). Condition-based maintenance conducts maintenance activities based on data gathered from condition-monitoring of equipment, as opposed to the traditional scheduled or breakdown maintenance [1]. Data is used to provide advance warning of failure so that equipment repair or replacement can be scheduled [2]. For deployable units, where the impact of a single equipment failure can severely limit mission effectiveness, CBM aims to minimize sudden failure when deployed and maximize the maintenance conducted during scheduled availabilities. 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 [3]. However, both the data and the understanding of the data are fundamental to any CBM program. Extracting actionable information from the raw data requires pattern analysis and an understanding of the underlying physics. One power monitoring option, nonintrusive load monitoring, is a low-cost solution for extracting device-level information from aggregate power data at a central location [4]. For a Nonintrusive Load Monitor (NILM) to be an effective CBM tool, both anomalous load behaviors and load signatures have to be detected and identified. Furthermore, the data should be presented as an intuitive decision aid for users.

This paper uses case-studies of faults observed with nonintrusive load monitoring aboard two United States Coast Guard cutters (USCGCs), SPENCER and ESCANABA, to present metrics and visualizations that have proven useful for CBM. First, the NILM installations aboard SPENCER and ESCANABA are described. Then some common faults in electromechanical systems are explained. Some useful metrics, or diagnostic indicators, for diagnosing those faults are presented. Finally, the paper discusses how to determine the proper fault warning levels and detect degrading equipment signatures, in order to use the described metrics for diagnosing fault conditions.

Nonintrusive Load Monitoring Onboard US Coast Guard Cutters

The USCGCs SPENCER and ESCANABA are 270 ft (82 m) medium endurance cutters (MECs) based in Boston, MA. The ships each maintain a 100-person crew, 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 present 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.

This article is based on research originally presented at AUTOTESTCON 2019 [17] and recognized as the Best Student Runner-Up Paper.

1094-6969/20/$25.00©2020IEEE

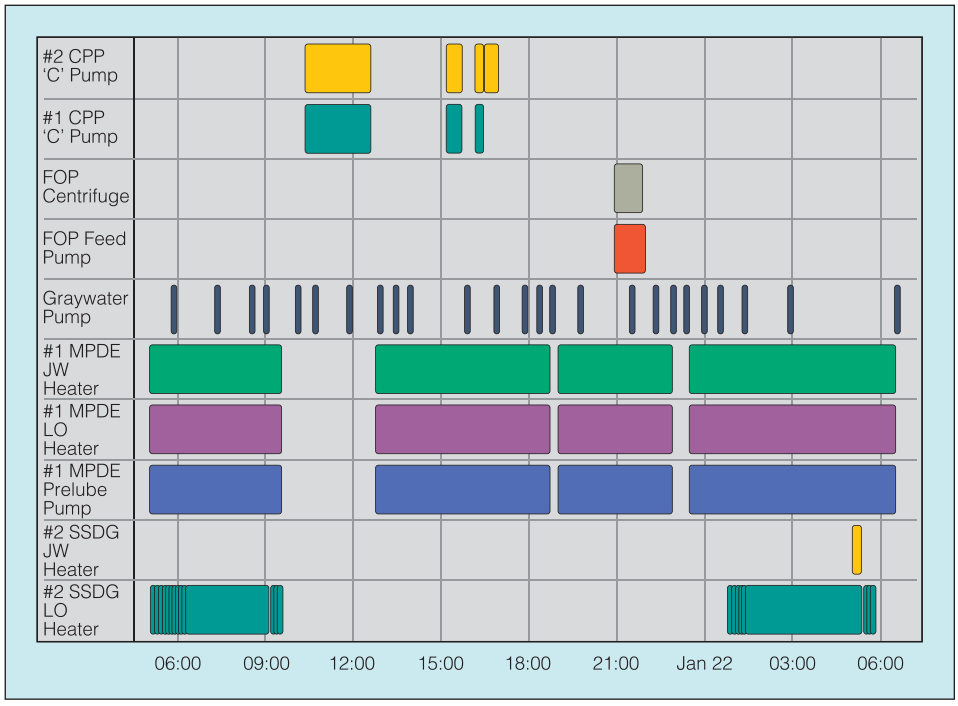

Fig. 1. The Timeline view of the NILM dashboard displays equipment status over a day of at sea operation. Colored blocks indicate periods where equipment is energized.

Fig. 2. NILM dashboard Metrics view indicating healthy graywater pump operation.

These sensor feedback systems are crucial to equipment operations, but they make equipment anomalies difficult for watchstanders to detect.

NILM systems are installed upstream of two 440 V sub-panels in the main engine room. These panels, STBD and PORT, 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 main propulsion diesel engines (MPDE) propelling the ship, as well as the two ship-service diesel generators (SSDG) providing power to the ship's microgrid. The NILM systems installed on these vessels serve two primary objectives: first, to identify equipment operating schedules to improve watchstander situational awareness; then, once the operating schedules have been accurately identified, to analyze the gathered data for CBM and fault detection to improve system operational availability. Information is presented to the crew onboard SPENCER through the on-site NILM dashboard [5].

A NILM system contains a NILM meter that uses a data acquisition unit (DAQ) to sample current and voltages at 8 kHz and then transmits the data via Ethernet to a host computer. 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 [6]. P and Q correspond to the envelopes of in-phase and quadrature current drawn relative to the load voltage [7]. 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 [8]. The operating schedule is communicated to the crew through the NILM dashboard "Timeline View," as shown in Fig. 1. 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. These metrics are displayed in Fig. 2 through the NILM dashboard "Metrics View" and are subsequently described.

Faults in Electromechanical Systems

Electrical monitoring through a NILM system can be used to monitor deviations from acceptable behavior or standard condition of equipment. The cause of these deviations are often "soft faults" or the gradual degradation of equipment performance. With the aid of automatic controllers and feedback control, system operability is maintained; however, this

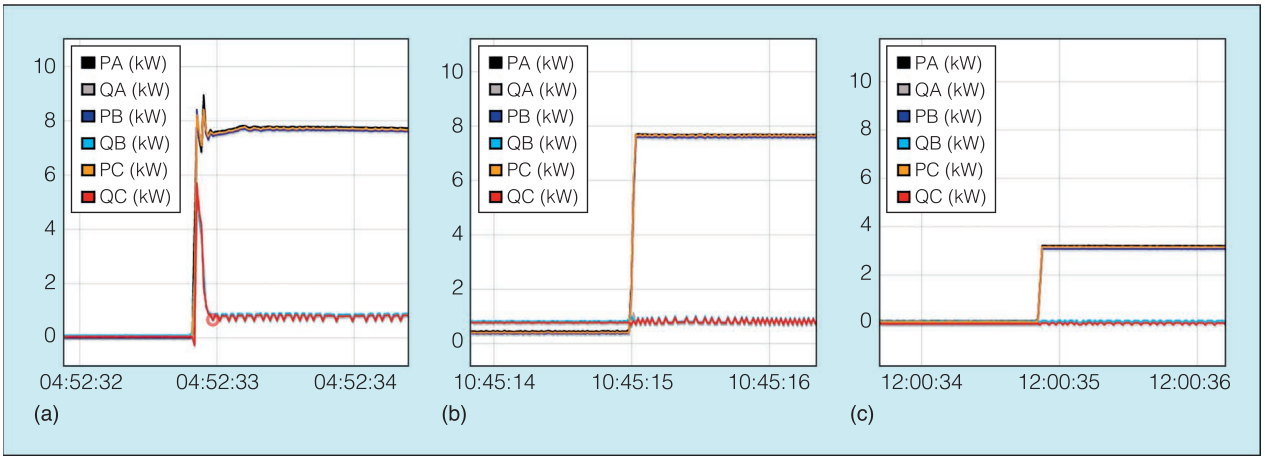

Fig. 3. *On* transient of MPDE: (a) pre-lube pump, lube oil heater, and jacket water heater, (b) lube oil heater and jacket water heater, and (c) jacket water heater [17]. Used with permission, (©IEEE, 2019).

often leads to increased wear and if left undetected, it eventually leads to a "hard fault," or the complete failure of a piece of equipment. Condition-based maintenance aims to detect soft faults, before they turn into hard faults. These faults are often invisible to watchstanders but clearly visible in the electric power readings.

Case Study: Main Propulsion Diesel Engine System

A main propulsion diesel engine (MPDE) jacket water (JW) heater failure from SPENCER and ESCANABA is presented here as an example of a soft fault failure that was invisible to watchstanders but clearly visible in the electrical readings. The MPDE JW heater consists of two 4.5 kW heating elements on either side of the engine block. The heaters appear to the NILM as a single electrical load, a 9 kW resistive load with 3 kW per phase. The MPDE JW heater is run by a controller and runs automatically based on 90° F and 120° F setpoints. The controller configuration results in the MPDE JW heater frequently turning *on* and *off* simultaneously with the MPDE lube oil (LO) heater and MPDE prelube (PL) pump. Fig. 3 shows the electrical transients for a healthy MPDE JW heater operating alone and in tandem with the other MPDE system loads. On both ships, the NILM detected a slow change in the heater's electrical signature, with the real power decreasing in a series of degradation events. During the same period, the heaters showed an increase in reactive power on different phases of power. Fig. 4 shows the power draw as the MPDE JW heater on ESCANABA port-side starts at healthy operation, then goes through degradation for about a minute. During this process, the MPDE JW heater decreases about 800 W on two phases of real power and splits in reactive power of about 450 VAR of opposite signs on two phases.

After this degradation event occurs, any time the MPDE JW heater turns *on* or *off*, it is operating in the degraded state with reduced heating capability. The MPDE JW heater went through several more degradation events similar to Fig. 4 over the course of observation. By plotting the real (P) and reactive (Q) power step change at each *on* or *off* event, the changes in the power draw after each degradation event is evident. Using ESCANABA port-side events from April 2018 to May 2019, Fig. 5

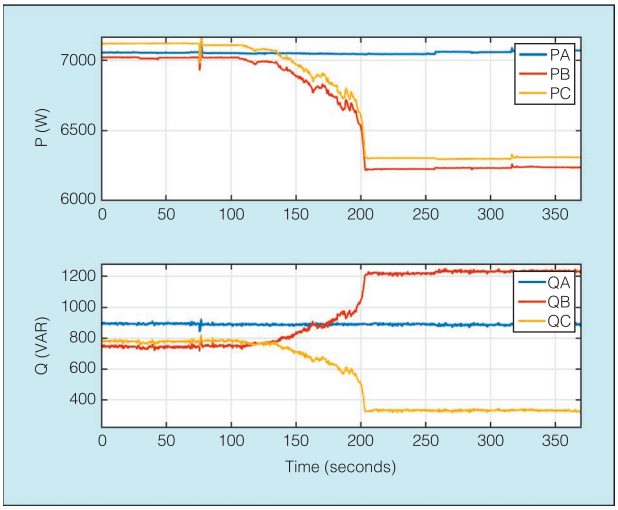

Fig. 4. Power draw as the MPDE JW heater degrades over the course of about a minute, decreasing about 800 W on two phases of real power and split of about 450 VAR in reactive power of opposite signs on two phases.

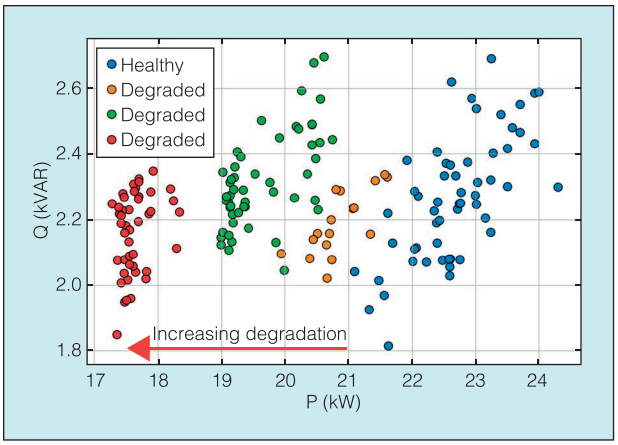

Fig. 5. The PQ space shows how the MPDE system cluster moves as the jacketwater heaters continue to degrade over time.

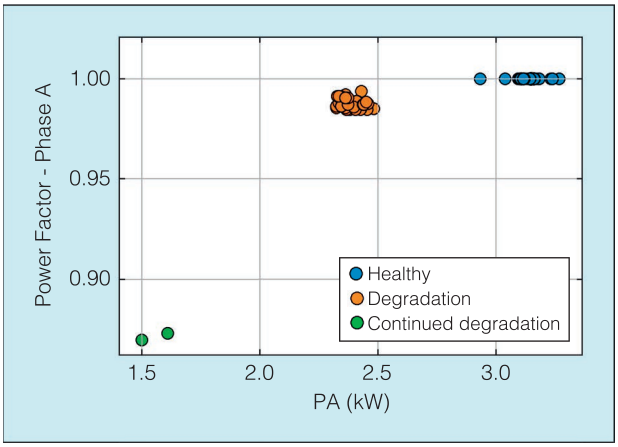

Fig. 6. The P-PF space shows how the MPDE JW cluster moves as the jacketwater heaters continue to degrade over time.

plots the PQ space for *on* events, as a sum of the three phases, of the MPDE JW heater in tandem with the LO heater and PL pump. Over this time period, the MPDE went through three degradation events, as indicated with the different clusters. The plot shows that with increasing degradation, the power draw continues to decrease. The power factor (PF) can also provide an indication of a fault. Normally a purely resistive load with a PF of 1, over the course of degradation the PF on various phases also decreases. This is shown in Fig. 6 for one phase of the SPENCER port-side heater as it degrades (using data from November 2016 to November 2018).

For a variety of reasons, this fault was nearly impossible to detect without the assistance of the NILM. Despite the holes in the heaters shown in Fig. 7, 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 to activate the load. After the fault condition was made visible by the NILM, the ships' crews were alerted and the heaters were subsequently removed and inspected. This revealed severe corrosion of the heating elements. Fig. 7 shows the physical condition of the heaters after their electrical signature decayed. Upon removal of the enclosure covers, the wiring was degraded and one of the heaters was lightly smoking. In addition to identifying a failed heating element, analysis detected a potential shock hazard and likely prevented an electrical fire in the engine room. There is currently no preventative maintenance 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.

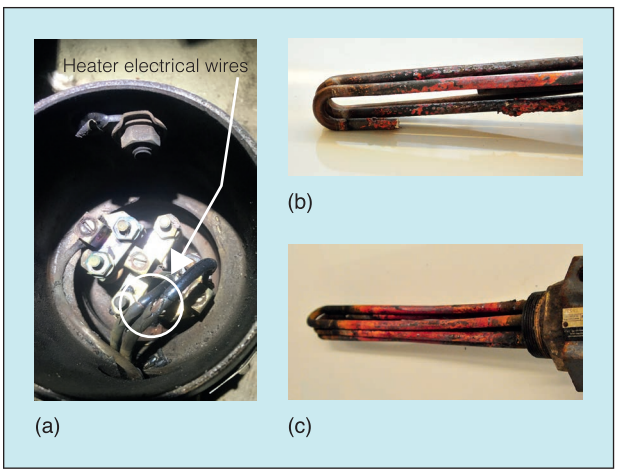

Fig. 7. (a) Bare wiring found in SPENCER port inboard jacket water heater. Jacket water heaters removed from engines after electrical degradation was observed, revealing severe corrosion: (b) from ESCANABA port inboard heater, removed April 2019 and (c) from SPENCER port inboard heater, removed December 2018.

Diagnostic Indicators for Fault Detection

Nonintrusive load monitoring records 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. An effective CBM program using a NILM should track changes in *equipment signature*, as demonstrated by the MPDE JW heater fault presented above, as well as changes in *equipment behavior*, demonstrated with the examples in [9].

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. The equipment in this work consists largely of pumps and heaters, each with a consistent steady-state power signature. *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. A change in *power* demand may indicate a worn motor bearing [10] or a change in *power factor* could be a sign of corroded heating elements, such as the MPDE JW heater example presented. Many of the heaters and pumps monitored by the NILM on the MECs are controlled by closed-loop automated systems such as tank-level sensors or thermostats. Thus, the *average run duration*, *total run time*, and *daily actuations* track equipment behavior and are useful for finding sensor and automation faults that might cause equipment to run too frequently or not enough. A broken tank level indicator or failed thermostatic controller can, for example, cause equipment to activate in repeated short-cycles or run for excessively long periods [9].

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 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.

Determining Fault Warning Levels

Condition-based maintenance parameters are communicated on the NILM dashboard via "green- yellow-red" diagnostic gauges, as shown in Fig. 2 [11]. 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 it a useful, actionable tool for the ship's crew. A variety of methods have been proposed to determine fault thresholds for industrial applications [12]. 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 [13].

Historical data collected by the NILM was used for SPC analysis. *Metrics* were calculated for the 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 [13]. The UCL and LCL are defined as:

$$UCL = \theta + k\sigma \tag{1}$$

$$LCL = \theta - k\sigma \tag{2}$$

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 [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" [13]. 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 [14]. Fig. 8 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 likelihood that the variable falls within that particular range of values. The quantile values displayed at the bottom of Fig. 8 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):

$$F(x) = Pr(X \le x) \tag{3}$$

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

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

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:

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

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. 8.

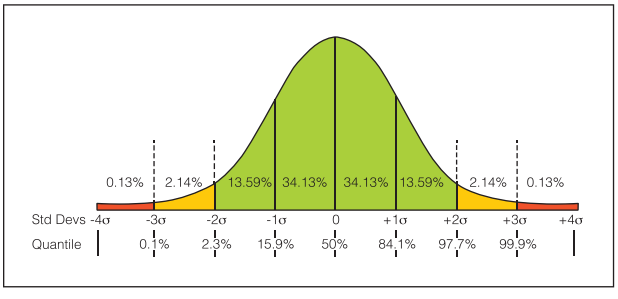

Fig. 8. Probability density function of a normal distribution showing progressive thresholds for fault detections. Colors correspond to the red, yellow, and green regions on the dashboard gauges [11]. Used with permission, (©MIT, 2019).

The ICDF function for a Weibull distribution is:

$$x = F^{-1}\left(p|\alpha,\beta\right) = \alpha\left(-ln\left(1-p\right)\right)^{\frac{1}{\beta}} \tag{6}$$

Therefore, the centerline can be found by setting p equal to 0.50 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 [11]. 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 distribution function will be used instead of the ICDF. The generalized inverse distribution function can be expressed as:

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

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 SPC method is visualized in Fig. 9 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 Fig. 9 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 [15]. The test rejects the null hypothesis at the 5% significance level. For the example shown in Fig. 9, 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 was used since it is applicable for discrete distributions [15]. 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. Depending on the PDF or PMF selected, the fault detection thresholds are set as shown in Fig. 8 or using a quantile function, as described in (6) and (7). This process was repeated to determine control limits for each of the monitored loads.

Detection of Changing Load Signatures

If a user knows that a piece of equipment should be operating, but the dashboard Metrics View shows the load has not been operational, this would alert the user that a fault is present. There are several reasons the *daily actuations* parameter could drop to zero, even if the load should be operating. First, it could be indicative of a broken sensor, such as a tank-level sensor or temperature sensor, resulting in the load not turning on even when it should be. Second, it could be indicative of a complete

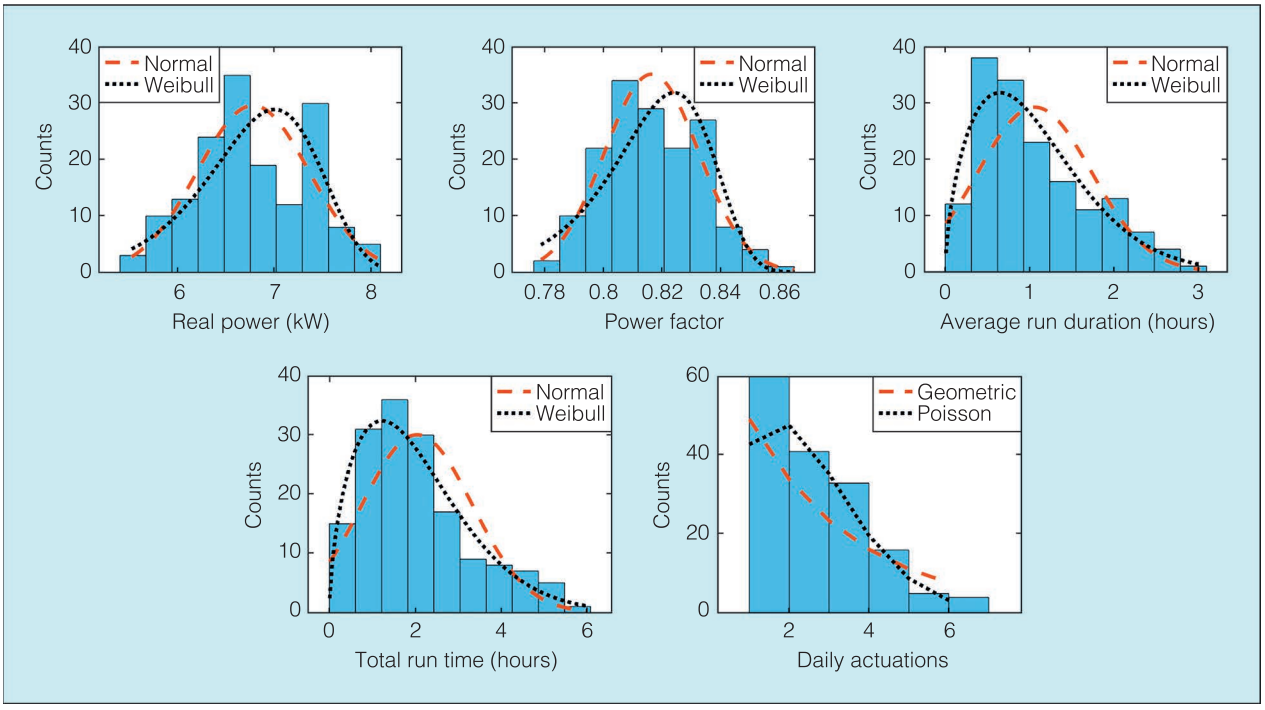

Fig. 9. Different probability density functions fit to each parameter of the controllable pitch propeller (CPP) pumps.

failure of a piece of equipment. Alternatively, it could be due to a degraded piece of equipment, resulting in a load signature that has changed to the point where load identification accuracy is reduced. However, as described, in order to help a user identify anomalous behavior, a CBM program using nonintrusive load monitoring should track changes not only in *equipment behavior*, but also *equipment signature*.

NILM identification algorithms are trained on healthy signatures, and it is not practical to train on all anomalous signatures because the changes depend on the root cause of the anomaly. Also, it is generally not clear to what extent a load signature could change before affecting load classification accuracy. Despite these challenges, knowledge of the equipment operation can aid in defining those boundaries and better identifying equipment, even in its degraded state. One way to identify degrading equipment is if it is part of a finite state machine (FSM) or an interdependent system. For example, in this work we were able to diagnose the MPDE JW heater fault condition even as its signature was changing, due to its simultaneous turn-on and turn-off with the other MPDE system loads and with knowledge of the MPDE system FSM behavior.

Next, knowledge of the actuation mechanism can aid in determining if changes in transient shape are due to natural variations or due to an underlying fault condition. When a load turns on it has an in-rush current, leading to a load transient that is often repeatable because the transient behavior of a load is related to the physical task the load is performing. However, because the in-rush current also depends on the time instant with respect to the voltage line-cycle that a load turns on, the variability in the turn-on transient also depends on the mechanism for turning-on [16]. Again, using the MPDE system as an example, the PL pump has an in-rush current that has a peak value that varies greatly, as seen in Fig. 10. Here, P_{peak} and Q_{peak} are the difference between the max power of the transient and the steady state power for real and reactive power, respectively. Because this load is activated with a mechanical switch, the load can turn on at any point in the voltage line-cycle. Thus, even with the variability in the transient peak for the MPDE PL pump, it is still operating in a healthy condition. Contrarily, if the load were actuated with a solid-state relay with zero-crossing detection, the variability in the transient peak would become significantly smaller, and a large variation in the transient peak would be a sign of a possible anomaly.

Although these are some examples of how anomalous load signatures can still be detected, it should be noted that it is still an open research question to resolve the balance between accurately detecting changing load signatures and not misclassifying loads.

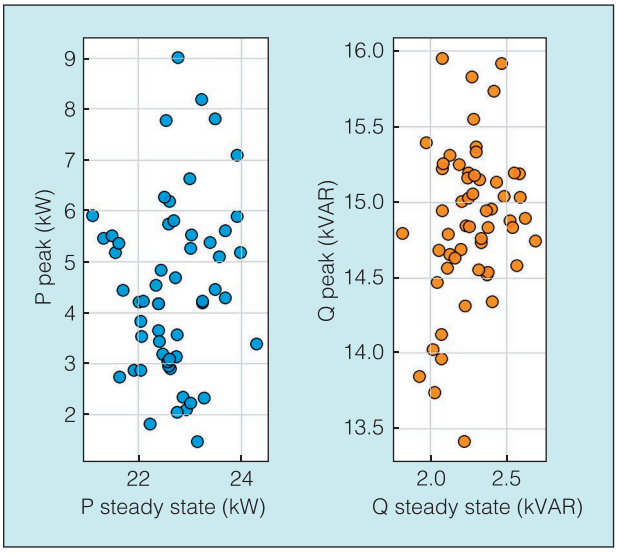

Fig. 10. Feature space of steady state power versus peak power shows the variation in peak power when the MPDE pre-lube pump turns on.

Conclusion and Future Work

As this paper has shown, a NILM has proven to be a realizable tool for CBM. However, there are still avenues for improvement. Methods for determining the boundaries in the load identification feature space would help determine to what extent a load signature could change before affecting load identification accuracy. For example, a method similar to the one implemented for creating *metrics* boundaries could be used to create green-yellow-red regions in the identification feature space, enabling more accurate fault detection. With accurate detection of faults and anomalous behavior prior to the ship's deployment, servicing and replacement of equipment can be better scheduled and operational availability can be maximized. Nonintrusive load monitoring represents a low-cost method for the rapid implementation of a CBM and fault detection program for electromechanical equipment. Analysis of USCGC loads show that through careful selection of metrics and simple statistical analysis, a 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] C. Chen, B. Zhang, and G. Vachtsevanos, "Prediction of machine health condition using neuro-fuzzy and bayesian algorithms," *IEEE Trans. Instrum. Meas.*, vol. 61, no. 2, pp. 297-306, Feb. 2012.

[2] H. M. Hashemian and W. C. Bean, "State-of-the-art predictive maintenance techniques," *IEEE Trans. Instrum. Meas.*, vol. 60, no. 10, pp. 3480-3492, Oct. 2011.

[3] X. Dai and Z. Gao, "From model, signal to knowledge: a data-driven perspective of fault detection and diagnosis," *IEEE Trans. Industrial Informatics*, vol. 9, no. 4, pp. 2226-2238, Nov. 2013.
[4] M. Figueiredo, B. Ribeiro, and A. de Almeida, "Electrical signal source separation via nonnegative tensor factorization using on site measurements in a smart home," *IEEE Trans. Instrum. Meas.*, vol. 63, no. 2, pp. 364-373, Feb. 2014.
[5] A. Aboulian, 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 Trans. Industrial Informatics*, vol. 15, no. 3, pp. 1405-1414, Jun. 2018.
[6] J. Paris, J. S. Donnal, Z. Remscrim, S. B. Leeb, and S. R. Shaw, "The sinefit spectral envelope preprocessor," *IEEE Sensors J.*, vol. 14, no. 12, pp. 4385-4394, 2014.
[7] C. Laughman, K. Lee, R. Cox, S. Shaw, S. Leeb, L. Norford, and P. Armstrong, "Power signature analysis," *IEEE Power and Energy Mag.*, vol. 1, no. 2, pp. 56-63, Mar. 2003.
[8] D. Green, S. Shaw, P. Lindahl, T. Kane, J. Donnal, and S. Leeb, "A multi-scale framework for nonintrusive load identification," *IEEE Trans. Industrial Informatics*, Jun. 2019.
[9] P. A. Lindahl, D. H. Green, G. Bredariol, A. Aboulian, 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] R. Beebe, "Estimate the increased power consumption caused by pump wear," *Pump Magazine*, vol. 58, pp. 20-27, 2008.
[11] T. J. Kane, "The NILM dashboard: shipboard automatic watchstanding and real-time fault detection using non-intrusive load monitoring," Master's thesis, Massachusetts Institute of Technology, 2019.
[12] S. Dash and V. Venkatasubramanium, "Challenges in the industrial applications of fault diagnostic systems," *Computers and Chemical Engineering*, vol. 24, pp. 785-791, 2000.
[13] D. C. Montgomery, *Introduction to Statistical Quality Control*. Hoboken, NJ, USA: John Wiley and Sons, Inc., 2009.
[14] W. A. Levinson, *Statistical Process Control for Real-World Applications*. Boca Raton, FL, USA: CRC Press, 2011.
[15] M. J. de Smith, *Statistical Analysis Handbook*. Edinburgh, Scotland, UK: The Winchelsea Press, Drumlin Security Ltd., 2018.
[16] M. N. Meziane, T. Picon, P. Ravier, G. Lamarque, J. Le Bunetel, and Y. Raingeaud, "A measurement system for creating datasets of on/off-controlled electrical loads," in *Proc. IEEE 16th Int. Environment and Electrical Eng. Conf. (EEEIC)*, pp. 1-5, Jun. 2016.
[17] D. Green, P. Lindahl, S. Leeb, T. Kane, S. Kidwell and J. Donnal, "Dashboard: Nonintrusive Electromechanical Fault Detection and Diagnostics," in *Proc. 2019 IEEE AUTOTESTCON*, pp. 1-9, 2019.

Daisy Green **(**dhgreen@mit.edu**)** is currently pursuing a Ph.D. degree in electrical engineering with the Research Laboratory of Electronics, Massachusetts Institute of Technology. She received a B.S. degree from the University of Hawaii at Manoa in 2015 and an M.S. degree from Massachusetts Institute of Technology in 2018, both in electrical engineering.

Thomas Kane is a Lieutenant in the US Coast Guard, stationed as the Engineering Officer onboard USCGC CAMPBELL. He received an M.S. degree in mechanical engineering at the Massachusetts Institute of Technology in 2019. He was previously stationed as Damage Control Assistant aboard USCGC MELLON and as a Port Engineer for the National Security Cutter (NSC) fleet.

Stephen Kidwell is a Lieutenant in the US Coast Guard pursuing an M.S. degree in naval architecture and marine engineering at the Massachusetts Institute of Technology. He served as a Student Engineer and eventually Assistance Engineer Officer onboard USCGC HAMILTON (WMSL 753) from 2015-2018.

Peter Lindahl currently works as an Engineering Consultant at Exponent in Watertown, Massachusetts. He was previously a Postdoctoral Associate at the Massachusetts Institute of Technology. He graduated with his Ph.D. degree in electrical and electronics engineering from Montana State University in 2013. His research interests include sensors and instrumentation for energy and power systems, renewable energy generation, and energy policy.

John Donnal is currently a Faculty Member of the U.S. Naval Academy, Annapolis, Maryland, in Weapons and Systems Engineering. He received the B.S. degree from Princeton University in 2007 and the M.S. and Ph.D. degrees from the Massachusetts Institute of Technology in 2013 and 2016, respectively, all in electrical engineering. His research interests include nonintrusive load monitoring synthesis, energy harvesting, and communication systems.

Steven B. Leeb has been a member of the Massachusetts Institute of Technology (MIT) Faculty, with the Department of Electrical Engineering and Computer Science, since 1993. He also holds a joint appointment with the Department of Mechanical Engineering at MIT. He received a Ph.D. degree from MIT in 1993. His research interests include the development of signal processing algorithms for energy and real-time control applications.