The NILM Dashboard: Shipboard Automatic Watchstanding and Real-Time Fault Detection using Non-intrusive Load Monitoring

by

Thomas John Kane
B.S., United States Coast Guard Academy (2012)

Submitted to the
Department of Mechanical Engineering
in partial fulfillment of the requirements for the degrees of
Master of Science in Mechanical Engineering
and
Master of Science in Naval Architecture and Marine Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2019

Signature redacted

Author .....
Department of Mechanical Engineering
May 10, 2019

Signature redacted

Certified by .....
Steven B. Leeb
Professor
Thesis Supervisor

Signature redacted

Certified by .....
Daisy H. Green
Doctoral Candidate
Thesis Supervisor

Signature redacted

Accepted by .....
Nicolas Hadjiconstantinou
Chair, Department Committee on Graduate Students

The NILM Dashboard: Shipboard Automatic Watchstanding and Real-Time Fault Detection using Non-intrusive Load Monitoring

by

Thomas John Kane

Submitted to the Department of Mechanical Engineering
on May 10, 2019, in partial fulfillment of the
requirements for the degrees of
Master of Science in Mechanical Engineering
and
Master of Science in Naval Architecture and Marine Engineering

Abstract

Non-intrusive Load Monitoring (NILM) measures power at a central point in an electrical network and disaggregates individual load schedules from the overall power stream. This thesis presents the NILM Dashboard, a data-analysis and user interface tool that provides real-time machinery monitoring and fault diagnostics using NILM data. The Dashboard was developed and deployed for use onboard US Coast Guard Cutters to act as an automatic watchstander and condition-based maintenance aid. The effectiveness of the system is demonstrated on power data collected from electrical panels in the ship's engine room. Case studies are used to evaluate the Dashboard's ability to detect fault conditions in electromechanical systems.

Thesis Supervisor: Steven B. Leeb
Title: Professor

Thesis Supervisor: Daisy H. Green
Title: Doctoral Candidate

Acknowledgments

Thank you to all the groups and sponsors whose financial support made this work possible: the U.S. Coast Guard, the Office of Naval Research NEPTUNE Program, The Grainger Foundation, the MITei-Exxonmobil collaboration, and the Cooperative Agreement between the Masdar Institute of Science and Technology and MIT. I would first and foremost like to thank Daisy Green, whose talent and patience made this thesis possible. Keep hacking. Thank you to Dr. Steven Leeb, for giving me the opportunity to join the group and for providing constant enthusiasm. Thank you to Andre Aboulian, Jennifer Switzer and John Donnal for developing the Dashboard program and NILM Manager interface, upon which this thesis is based. Thank you to my fellow LEES-GEM students, Lukasz, Dayo, Steve K., Erik, Eric, and Manuel, for providing some welcome distractions in the lab. Thank you to the many U.S. Coast Guard personnel who assisted us with this research, especially LCDR Nicholas Galati, LT Devin Quinn, CWO Daniel Doherty, and EMC Jason Sardinas. Your substantial knowledge and dedication to duty were on display during our many visits to the cutters. A special thank you to my parents, for your constant support and for making me the student I am today. Thank you to Ashley, for being there when I needed to vent and for mostly staying awake when I described this thesis to you. Thank you to my great group of friends here in Boston, especially Peter, Roland, Jarrod, and Maddie, for helping make the last two years so much fun.

Contents

List of Figures

List of Tables

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 1

Executive Summary

Machinery watchstanders in the modern Coast Guard and other maritime fleets have vastly different experiences depending on the asset. On legacy ships, such as Medium-Endurance Cutters (WMECs) or Patrol Boats (WPBs), 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. On modernized assets, such as the National Security Cutter (WMSL), sensors bring data from nearly every pump, motor and valve to a single screen. Watchstanders can observe and control almost any system using a single streamlined interface. Each mode of machinery watchstanding brings a unique set of challenges.

For legacy assets, adding increased automation and sensing capability can reduce the workload and improve the situational awareness of the crew. However, adding significant monitoring capability within schedule and budgetary contraints is challenging. Legacy machinery plants are not easily integrated into MCMS systems and these projects are unattractive for vessels at or past their mid-life.

For modernized assets with centralized MCMS, the promise of big data and predictive analytics remains only a promise. MCMS and other automation systems have allowed the use of a minimally-manned model to drive down life cycle costs. However, crews of these vessels are finding maintenance requirements are not reduced by the collection of endless logs of machinery operational data. Crews, especially minimally-manned crews, cannot be expected to parse vast data sets to detect anamolies. Condition-based maintenance (CBM) requires a reliable baseline that is often unclear or unknown, especially on new assets. Finally, analysis of machinery data is normally provided by shoreside support after deployment, often too late to avert failures in mission critical equipment while underway.

Non-intrusive load monitoring (NILM) is a technology that can offer assistance to watchstanders on vessels at all stages of automation. The NILM system described in this report

has the ability to:

- Easily integrate into existing machinery plants to create a centralized machinery monitoring point.
- Analyze collected data in real-time and provide *actionable* equipment diagnostics and CBM recommendations.

NILM works by monitoring the current stream at a centralized point in the electrical system, for example "upstream" from a switchboard or sub-panel. By observing the electrical transients at this point, NILM can identify changes in electrical power as individual pieces of equipment and determine the operating schedule of downstream loads. The unique electrical transient for a specific load can also be analyzed to assess the equipment's health. From a single sensing point, NILM can act as both a machinery operation tracker and a condition-based maintenance tool for numerous pieces of electromechanical equipment.

1.1 Non-Intrusive Load Monitoring on WMEC-270

Non-Intrusive Load Monitors are currently installed onboard USCGC SPENCER and ESCANABA, two WMEC-270 Famous class cutters based in Boston. The machinery plants of these legacy cutters lack a fully integrated MCMS, but still contain many closed-loop, automated systems. Closed-loop systems are actuated by sensor feedback, such as a tank-level indicators or temperature sensors. These systems are crucial to machinery operation, but also make faults difficult for watchstanders to detect. The automated control will continue to run the system even after a fault occurs. For instance, a vacuum leak in the sewage system will cause the pumps to run much more frequently, but the system still draws vacuum and the toilets continue to flush. Faults of this nature are nearly invisible to watchstanders, but are clearly present in electric power readings and thus can be detected by NILM. The installations on the WMEC-270s serve two primary objectives; first, use NILM 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.

To maximize NILM's utility for the crew, NILM current sensors were placed upstream of two 440V sub-panels in the main engine room. These panels (3-117-1/2) supply a variety of different mission-critical systems, listed in table 1.1 below. These panels also supply the loads that support the Main Propulsion Diesel Engines (MPDE), Ship's Service Diesel Generators (SSDG) and Controllable Pitch Propellers (CPP). Figure 1-2 shows how these

two monitoring points can provide insight into all major systems in the WMEC-270 engine room. The NILM installation is located directly above the monitored sub-panels (Fig. 1-1) and has successfully gathered data over the course of multiple deployments on both cutters.

Table 1.1: Electrical Loads Monitored by NILM

Port Sub-panel (3-117-2)	**Starboard Sub-panel (3-117-1)**
#2 MPDE LO Heater	#1 MPDE LO Heater
#2 MPDE J/W Heaters	#1 MPDE J/W Heaters
#2 MPDE PreLube Pump	#1 MPDE PreLube Pump
#2 SSDG LO Heater	#1 SSDG LO Heater
#2 SSDG J/W Heater	#1 SSDG J/W Heater
#2 CPP "C" Pump	#1 CPP "C" Pump
Inport S/W Pump	Shaft Turning Gear Motor
Fuel Oil Purifier	Oily Water Separator
Graywater Pumps	Bilge and Ballast Pump

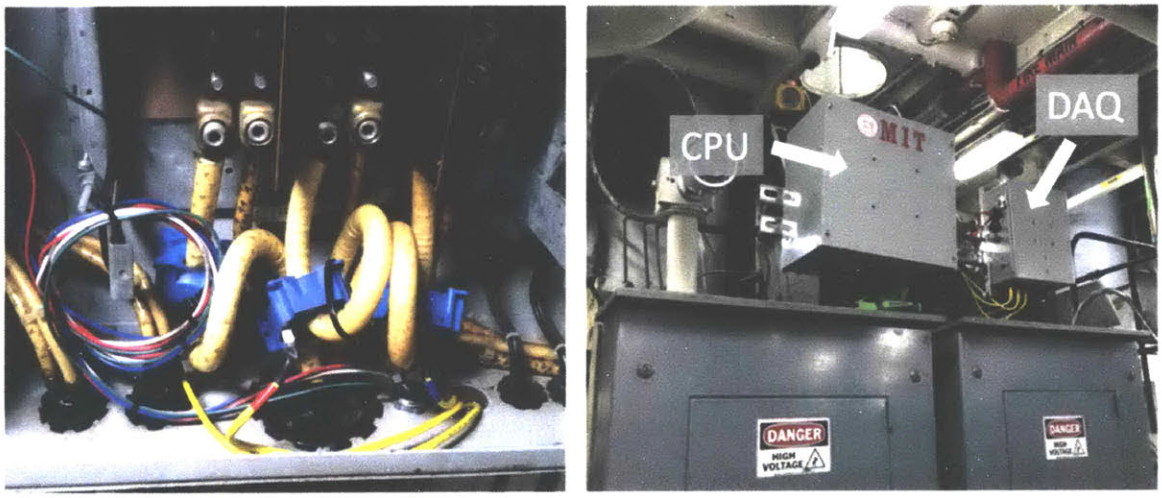

Figure 1-1: Left: NILM current sensors within a ship sub-panel. Right: NILM data acquisition (DAQ) and processing (CPU) units in SPENCER engine room.

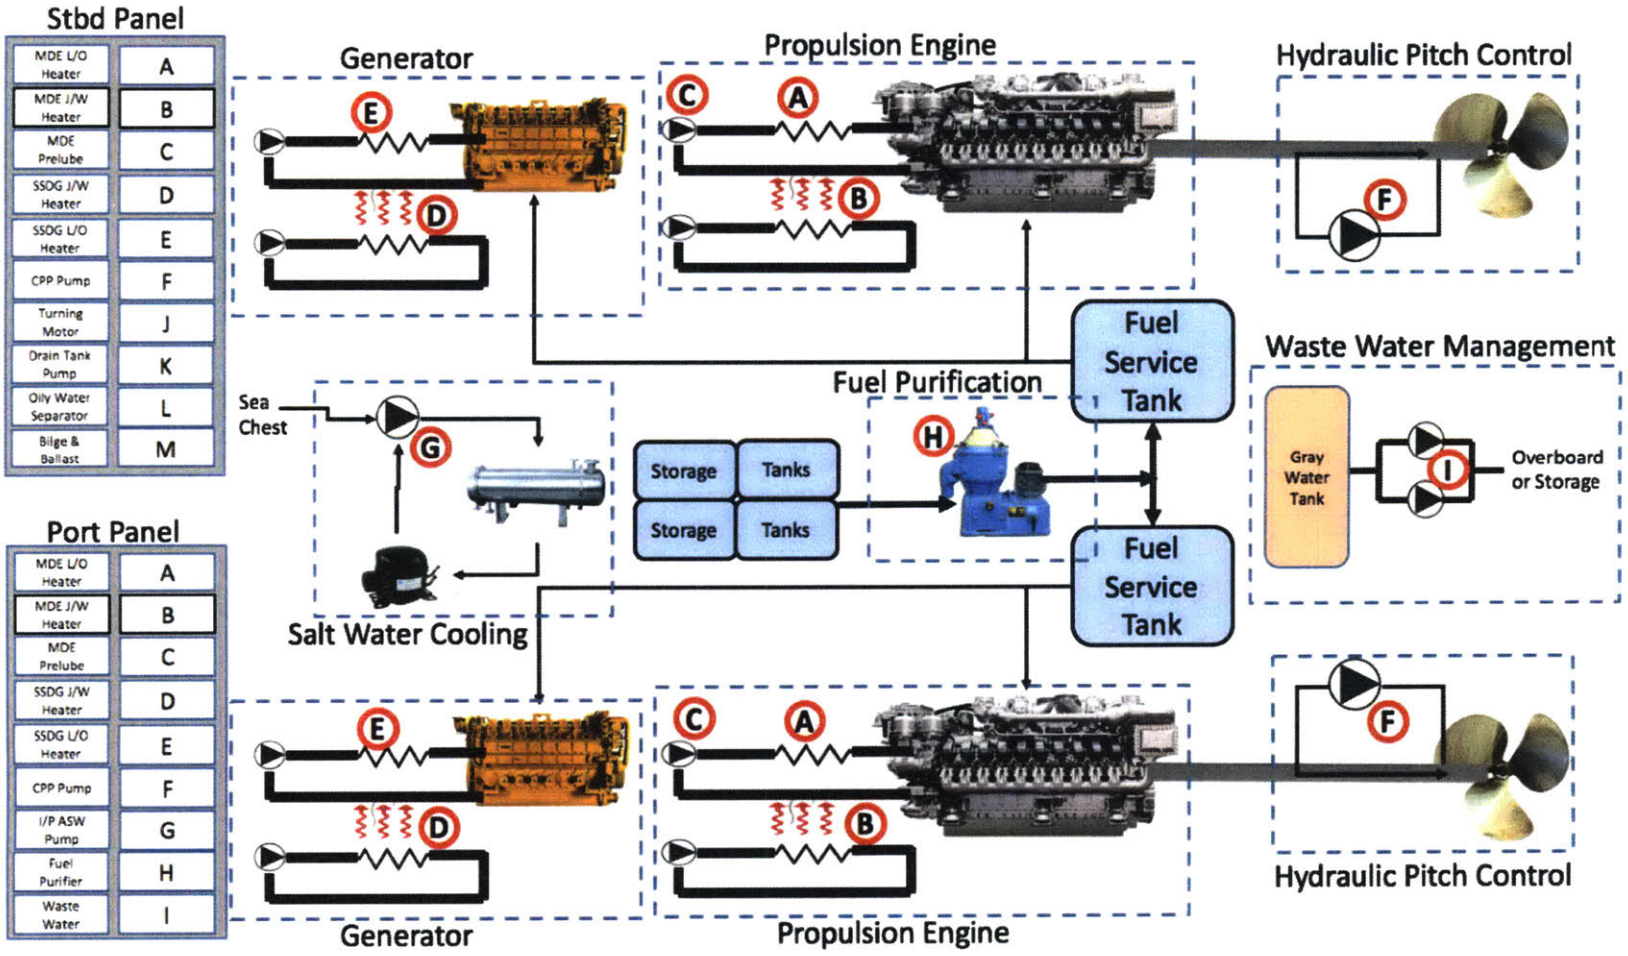

Figure 1-2: Engine room layout and monitored sub-panels [1]

1.2 NILM Dashboard

Since the NILM system has been deployed, the standard practice has been for MIT researchers to retrieve power data after each patrol, identify load events, and look for fault conditions within the system. However, the benefits of NILM cannot be fully realized unless the data is available to the crew in real-time. To achieve this, MIT deployed the NILM Dashboard onboard SPENCER in 2018. From the Dashboard, crew members can view equipment operating schedules and equipment diagnostics in an intuitive interface. Rather than simply adding a new set of data for the operators and maintainers, the Dashboard provides information that is immediately actionable for the crew. The installation on SPENCER can be seen in Fig. 1-3.

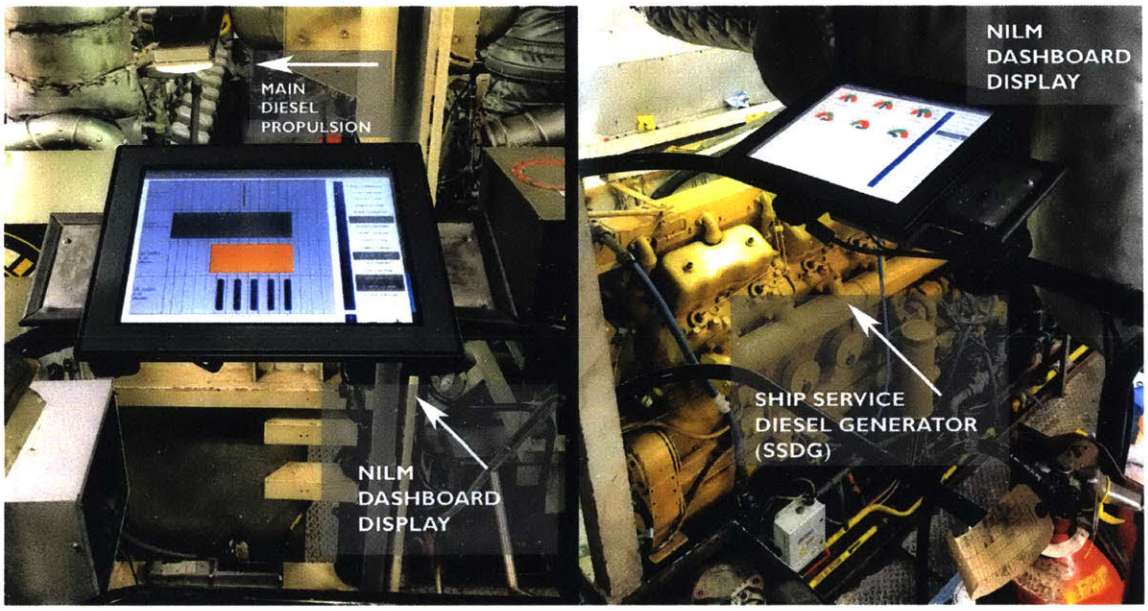

Figure 1-3: NILM Dashboard installation onboard USCGC SPENCER

The Dashboard has multiple different interfaces, each serving a unique objective. To improve situational awareness, the "Timeline" view displays the status and the recent operation of all the monitored loads. In addition, the Timeline can display the status of the SSDG and MPDE, based on the activation of the associated pumps and heaters. An example of the Timeline view from a busy day of underway operations can be seen in Fig. 1-4. With this data available, a crew member can quickly deduce the recent operating schedule of any of the 19 monitored loads.

The "Metrics" and "Historic" views serve as the primary touchpoint for condition-based maintenance. The Metrics view (Fig. 1-5) displays the health on familiar "green-amber-red"

gauges that cover five categories. The power and power factor gauges can detect material degradation of equipment such as mechanical wear and corrosion. The run duration, run time, and daily actuation indicators are used to find sensor and automation faults that might cause equipment to run too frequently or not enough. A combination of equipment nameplate data and the equipment behavior from past patrols were used to determine a baseline (i.e. the green regions on each gauge). Once a potential fault has been identified, the crew can use the Historic View (Fig. 1-6) to examine equipment health trends. This view simply displays the readings from the gauges from past days, and can inform the watchstander if the present state is an anomaly or part of a larger pattern. Together these interfaces can serve to direct watchstanders to previously invisible faults within the system.

Figure 1-4: The NILM Dashboard displays equipment status over a day of at sea operation. Colored blocks indicate periods where equipment is energized.

Figure 1-5: NILM Dashboard metrics indicating healthy gray water pump operation.

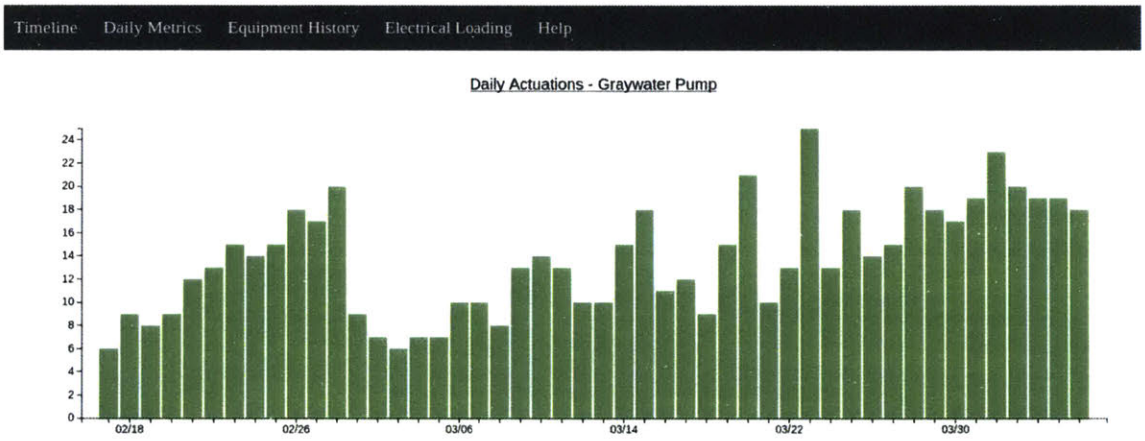

Figure 1-6: NILM Dashboard Historic view showing health gray water run frequency.

1.3 Results

Recent failures of MPDE jacket water heaters on SPENCER and ESCANABA showcase the utility of a NILM system. Each MPDE has two 4500W immersion heaters on each block that actuate when the water reaches a set point. NILM data showed that real power draw from these heaters was slowly degrading over time. This behavior was observed on all four engines on both cutters. After being alerted to the issue by MIT researchers, the ships' crew removed the heaters for inspection and discovered significant decay. There were large holes

in the heating elements, rendering them completely ineffective. Furthermore, the damage to the heaters was allowing jacket water to enter the heaters' electrical enclosure. Upon inspection of one enclosure, bare wire was found and the connections were lightly smoking. In addition to identifying a failed heating element, NILM analysis detected a potential shock hazard and may have prevented a Class Charlie fire in the engine room. A memorandum detailing this failure was submitted to USCG MECPL (Medium Endurance Cutter Product Line) in January 2019. The failed heaters are shown in Figures 1-7 and 1-8.

Figure 1-7: Top: New MPDE jacket water heater. Bottom: ESCANABA jacket water heater after removal.

Figure 1-8: Left: Bare wiring found in SPENCER jacket water heater. Right: Holes (circled) in the heating element that allowed water to enter electrical enclosure.

This fault was nearly impossible detect without the assistance of NILM. Despite the holes in the heater, 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 automated temperature sensors continued to activate the heaters at regular but only slightly more frequent intervals. There is currently no periodic maintenance action that prompts the crew to check the heater enclosures for damage or circuit continuity. Without closely monitoring the time it takes for engine jacket water to reach its temperature set points, there was little chance of crew members detecting this fault. The jacket water heater case is just one of several failures detected by NILM, covered in detail in chapter 4.

NILM represents a cost-effective tool that can reduce crew workload and improve machinery reliability on any platform. NILM systems have a small footprint and can be rapidly added or removed from systems that need additional monitoring. Vessels across the fleet could benefit immediately from the addition of NILM on mission critical systems. The remainder of this report describes the following:

- Chapter 2: Equipment monitoring on the WMEC-270
- Chapter 3: The use of NILM for automated watchstanding
- Chapter 4: The use of NILM for fault detection

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 2

Equipment Monitoring on the WMEC-270

The Non-Intrusive Load Monitoring Dashboard described in the previous chapter is the result of a power monitoring study that has been ongoing on USCGC SPENCER and ESCANABA since 2015. Developing a useful tool for machinery monitoring and condition-based maintenance using NILM is a multi-step process, illustrated in this study. Sensing hardware acquires raw data that is stored for review by researchers [1–3]. Fingerprint transient power signatures are identified by training an algorithm that automates time-intensive data correlation. Human annotation is typically still required to develop the best set of of fingerprint examples in this work, although automation of the process is an area of active research [4]. Load identification turns the stream of power data into an equipment event log. Finally, the equipment event log must be analyzed to determine what electrical signatures and equipment behavior are healthy, and what constitutes a fault condition. This information must then be presented to the machinery operators and maintainers in a user-friendly interface. Our experimental interface for this work is called the NILM Dashboard. This chapter briefly describes each step in the process, and details the healthy operation of a sample of monitored equipment. Diagnostic and fault detection applications will be discussed further in Chapter 4.

2.1 Data Acquisition and Storage

Non-intrusive metering hardware is located in the SPENCER and ESCANABA electrical systems as shown in Figure 2-1. There are two types of sensors deployed; contact and non-contact. The contact meters (shown in Figure 2-2) are deployed on two electrical subpanels that supply a variety of mission-critical loads. Through the load disaggregation

techniques described below, the contact meters are used to determine the operating schedules of equipment served from the sub-panels. Figure 1-2 shows how these two monitoring points can provide insight into all major systems in the WMEC-270 engine room. The non-contact meters consist of multiple sensors placed on the outside of a cable bundle and do not require ohmic contact for voltage detection (see Figure 2-3). These meters are used to measure the ship's total energy consumption and are placed on the four feeders that supply the main switchboard from onboard generators and shore power connections.

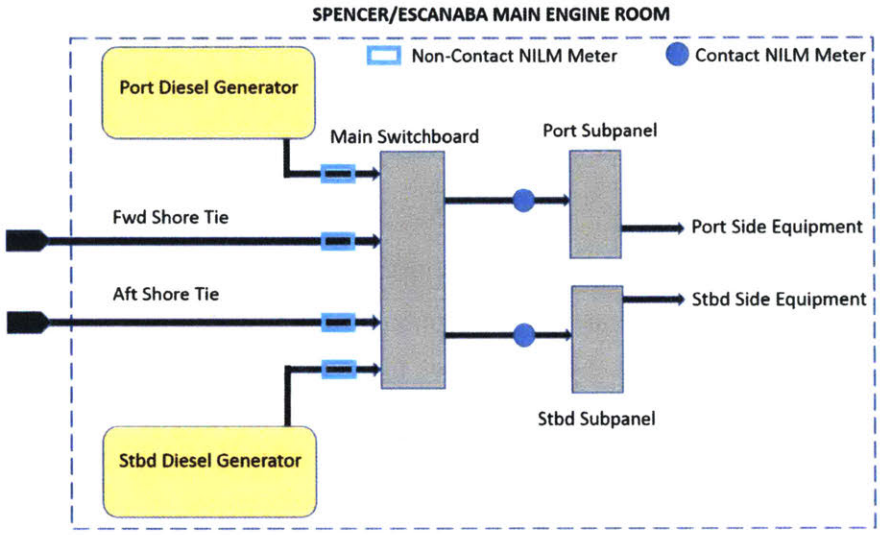

Figure 2-1: NILM metering hardware within WMEC-270 engine room.

The meters record current and voltage at a sampling frequency of 8 kHz for contact meters and 3 kHz for non-contact meters. These high sampling rates are necessary for capturing transient shapes as loads change state [5] and the higher harmonic content of non-linear loads [6]. To reduce data volume while maintaining the relevant information, these measurements are pre-processed in an onboard data acquisition unit into harmonic power envelopes using the Sinefit algorithm [7]. This process compresses the high-rate current and voltage data into real, reactive, and harmonic power streams at the system line frequency (60 Hz for the cutters). The processed data is streamed to an onboard computer via ethernet cable and stored in a high-speed time-series database known as NilmDB [8], from which the data can be viewed and analyzed in a web-based interface known as NILM Manager [3]. The flow of data from sensor to NILM Manager follows the same path regardless of the type of sensor used, and is shown in Figure 2-4. Note that NILM Manager is an interface for technical analysis of the power data and not intended for use by operators, unlike the Dashboard interface described in Chapter 1.

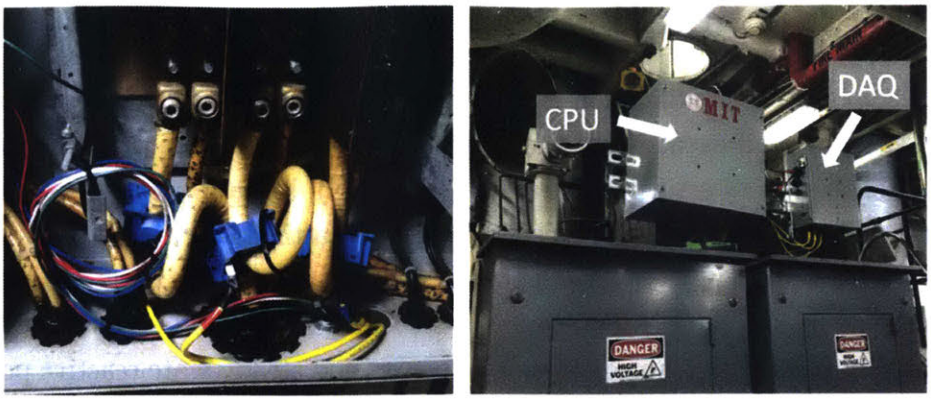

Figure 2-2: Left: Contact current sensors (blue) within a ship subpanel. Right: Contact data acquisition (DAQ) and processing (CPU) units in SPENCER engine room.

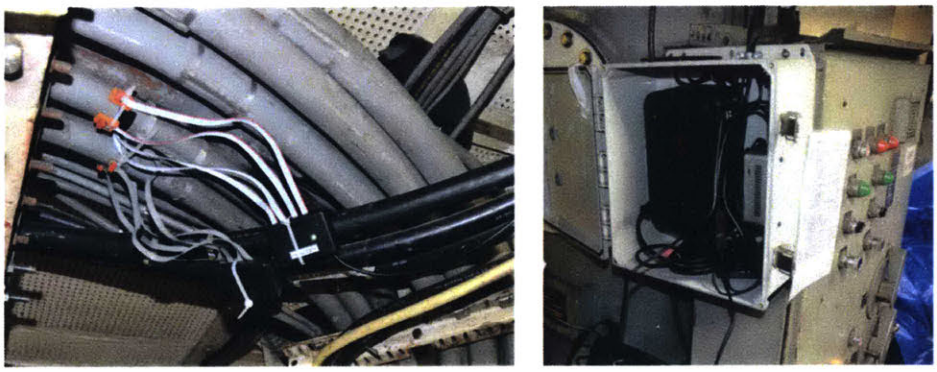

Figure 2-3: Left: Non-contact sensors placed on switchboard feeder cables on USCGC SPENCER. Right: Non-Contact DAQ and CPU, mounted near aft graywater system on SPENCER.

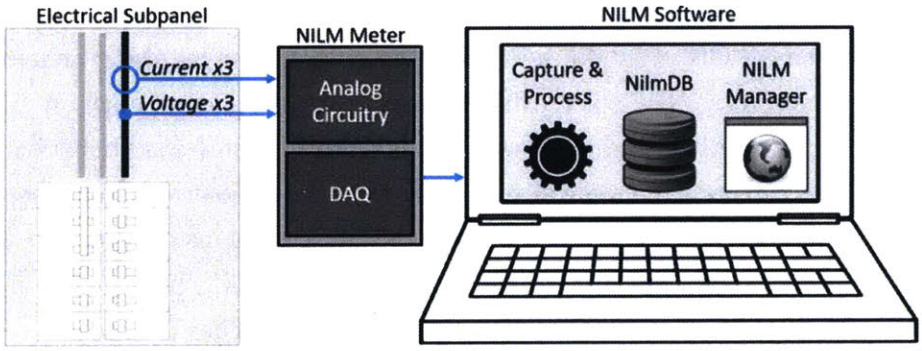

Figure 2-4: Depiction of data flow in the NILM system. Voltage and current readings are processed through a DAQ unit and stored on a CPU. Data flow is the same for contact and non-contact sensors [3].

2.2 Load Identification

The crucial step of load identification is conducted using the NILM Manager interface. Accurately identifying load events from the power stream can be accomplished by many different algorithms, such as artificial neural networks (NN), k-nearest neighbors (k-NN), and correlation-based algorithms [8, 9]. This study uses an artificial neural network approach. When a step change in real or reactive power is detected, a set of features are extracted to be used as the input into the NN. For example, the features extracted for an ON event might include transient peak power, steady-state real power, and steady-state reactive power. Successful load identification transforms a stream of electrical power into a time-series of discrete equipment events. Figure 2-5 shows a real power transient event and the associated features. Appendix B describes the load detection and identification process in greater detail.

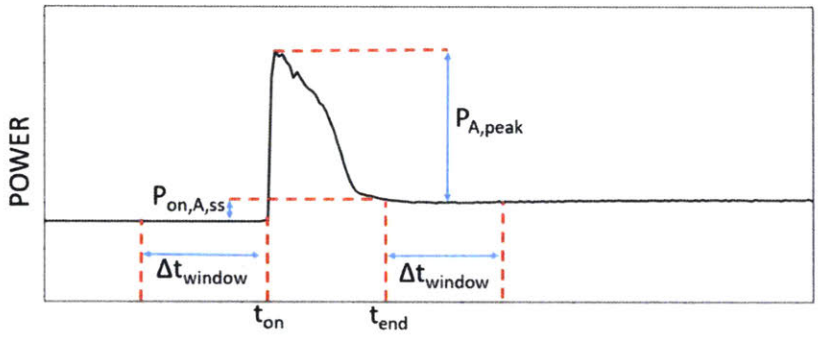

Figure 2-5: Transient ON event for a single phase of real power. $P_{on,x,peak}$ and $P_{on,x,ss}$ represent the peak power and steady state power features, respectively.

Successful load identification allows NILM to act as a machinery monitoring tool. The NILM Dashboard "Timeline view", shown in Figure 1-4 for a busy period of activity on CGC SPENCER, visualizes the disaggregated operating schedules for watchstanders on the WMEC-270. When an ON event for is detected, a colored bar for the corresponding load is generated on the screen. The colored bar remains on the Timeline until an OFF event is detected. The Timeline gives machinery watchstanders insight into load behavior that would be very difficult to discern on a typical round of the engine room. The interdependence of load events also allows NILM to infer the status of larger systems on the ship, such as the MPDEs and SSDGs. Load events can be used to generate a "Machinery Status" screen for the WMEC-270 engine room, shown below in Figure 2-6. The Timeline and Machinery Status screen are described in detail in Chapter 3.

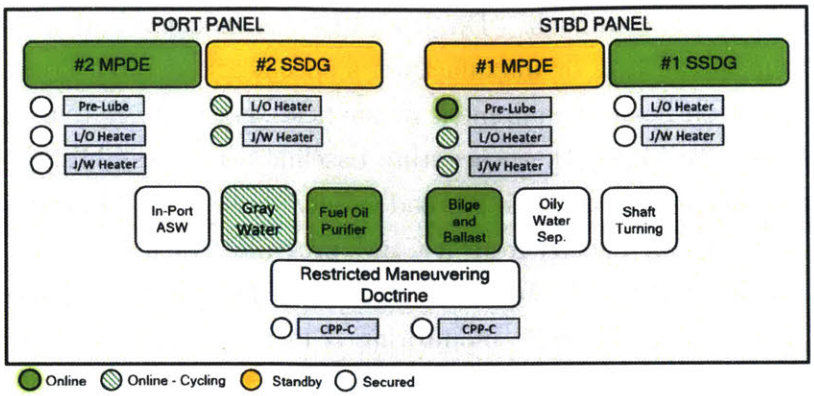

Figure 2-6: Visualization of machinery status in 270-WMEC engine room, determined using electrical data from the two monitored sub-panels.

2.3 Data Analysis for Condition-Based Maintenance

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, including the U.S. Coast Guard, 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 [10]. Preventative steps are only taken when there is evidence of degraded equipment performance. With a CBM program, data is used to provide advance warning of failure so that equipment replacement can be scheduled [10]. There is an even greater incentive to implement CBM onboard ships, where the impact of a single equipment failure at sea can severely impact the effectiveness of the entire asset. The logistics of acquiring a replacement part while far from home port further multiplies costs and asset downtime. An effective shipboard CBM program maximizes the maintenance conducted during scheduled inport periods and minimizes sudden failures when underway.

There are clearly advantages to implementing a condition-based maintenance program, but several significant challenges are present. First, establishing what constitutes healthy

equipment behavior can be difficult. Due to variations in environment and application, data gathered from lab testing or manufacturer's specifications often does not match the parameters measured on healthy equipment in the field. For some systems, such as the US Navy's LM-2500 Gas Turbine [11], an accurate baseline can be established by continually monitoring thousands of systems across a broad variety of use cases. However, many systems and organizations that can benefit from a CBM program will not have access to such a complete operational data set. Furthermore, 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 [12]. Reliable failure data could also be obtained by setting up a test facility to run equipment to failure in a variety of modes, but this is an unreasonable undertaking for most organizations.

Despite these challenges, a useful condition-based maintenance program *can* be implemented without a known baseline and complete set of failure data. Instead of using a pre-established healthy condition as baseline, the current condition of a piece of equipment serves as the starting point. Condition-monitoring can observe deviations in a parameter from the initial condition. These deviations may be the result of a "soft fault", where equipment continues to run, but with increased wear and energy consumption that eventually leads to failure. A soft fault warning can be given when a parameter has reached a threshold deviation from the initial condition. Because failure data is not available, the warning does not mandate that equipment must be repaired or replaced, but serves only to draw attention to potential failures. These warnings are highly valuable because soft faults are often invisible without condition-monitoring in place. Over time, as failures occur and more data is gathered, the fault thresholds can be adjusted to improve the accuracy of the program. By choosing the appropriate parameters and setting reasonable warning thresholds, a CBM program can be implemented without a large data set or up-front investment. Non-intrusive Load Monitoring in many respects represents the ideal way to develop a low-cost condition-based maintenance aid, with a small sensor footprint and easily-collated data sets. The success of NILM as a condition-based maintenance tool is documented in the case studies presented in Chapter 4.

2.3.1 Choosing CBM Criterion

The power data gathered through Non-intrusive load monitoring can be used to establish a shipboard CBM program. 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 (kilowatts) draw.
- Power Factor: The average ratio of real power to apparent power.
- Run Duration: Time between activation and shutdown.
- Runs Per Day: Number of discrete operations per day.
- Total Run Time: Total time the equipment is online over a 24-hour period.

These parameters work well for the equipment monitored onboard the WMEC-270, but are not an inclusive list of indicators for fault diagnostics with NILM. 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 Run Duration track the equipment signature, and Run Duration, Runs Per Day and Total Run Time track equipment behavior. The selected parameters were tailored for the equipment served by the monitored sub-panels. For example, the equipment sample in this work consists largely of pumps and heaters (see Table 1.1), each with a consistent steady-state power signature. Therefore, tracking the steady-state real power and steady-state power factor can detect equipment wear or changes in operational behavior. A change in real power demand may indicate a worn motor bearing [13] or a change in power factor could be a sign of heater corrosion (see Chapter 4). 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 [6].

Many of the heaters and pumps monitored by NILM on the WMEC-270 are controlled by closed-loop automated systems such as tank-level sensors or thermostats. Therefore, it is important to track the Run Duration, Runs Per Day and Total Run Time to detect broken sensing systems. A fouled tank level indicator or failed thermostatic controller can cause equipment to activate in repeated short-cycles or run for excessively long periods [14].

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 Runs Per Day 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. If the NILM does not detect a gray water pump activation for an extended period of time, this could be the indication of a fault. A 24-hour period was selected to observe run frequency for this study, but this parameter can be tailored to the equipment being monitored to more rapidly detect a fault condition, i.e. "Runs Per Hour."

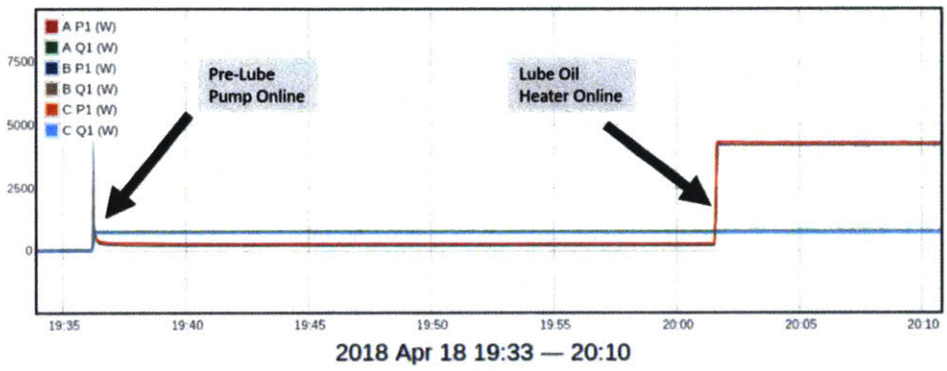

Figure 2-7: Normal sequence of operations when a MPDE engine secures. Pre-lube pump activates and then the lube oil heater activates shortly after. The presence of one load without the other can be used as a fault warning.

The interdependence of equipment can also be used diagnose fault conditions, and was useful in locating the MPDE jacket water heater fault detailed in Chapter 4. Finite State Machine (FSM) loads [15] have multiple stages of operation that must occur in the same order during normal operations. For instance, the MPDE lube oil heater should not be run unless the pre-lube pump is already online. Figure 2-7 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. Future iterations of the NILM Dashboard could add this capability to improve watchstander awareness of interdependent loads and FSMs.

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* of each parameter over 24 hours (see Appendix B). The 24-hour average serves to help prevent the gauges from falsely displaying an alarm as the result of a brief anomaly. The 24-hour averaging period can easily be adjusted for different applications where loads activate less or more frequently or tighter controls are required.

2.3.2 Determining Fault Warning Levels

Condition-based maintenance parameters are communicated on the NILM Dashboard via the "Green-Amber-Red" diagnostic gauges on "Metrics" interface (Figure 2-8 below). The amber region can be considered analogous to a trouble warning, whereas 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 [12]. For this study, a statistical process control (SPC) method is used. Effective SPC attempts to differentiate between natural variations and variations due to process failure [16].

Historical data collected by NILM was used for SPC analysis. The metrics generation script (see Appendix B) was run retroactively on NILM data collected from the SPENCER and ESCANABA sub-panels dating back to 2016. Readings were generated for each of the five parameters for every ten minutes of power data available. Equipment nameplate information and the ship's logs were used to exclude any data that may be 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 determine exactly how much deviation is normal and when a deviation should trigger a fault warning.

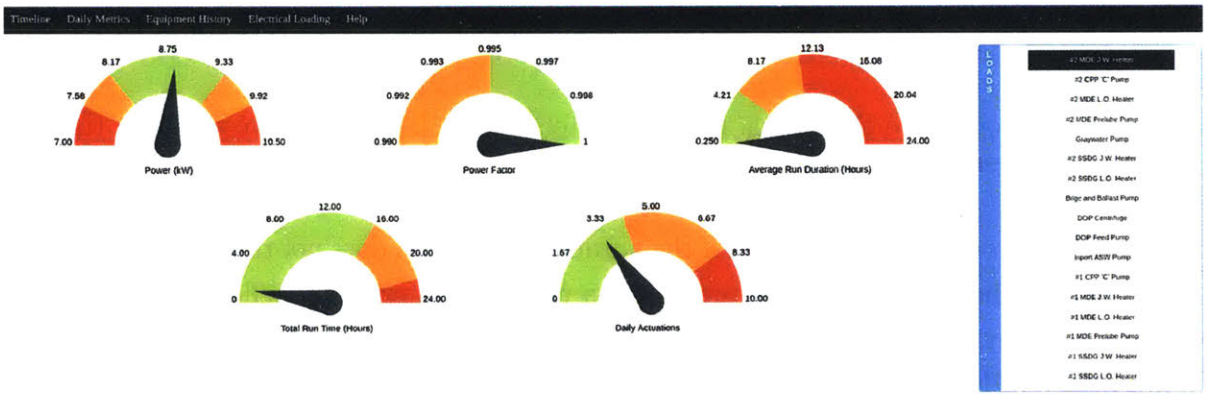

Figure 2-8: Dashboard gauges for jacket water heater showing a healthy condition.

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 control limits. These control limits can be thought of as the red regions of the gauges. For data with a normal distribution, SPC uses the arithmetic mean (θ) of the parameter as the centerline [16]. The UCL and LCL are defined as follows:

$$UCL = \theta + k\sigma \tag{2.1}$$

$$LCL = \theta - k\sigma \tag{2.2}$$

where σ is the standard deviation and k 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 [16,17]. This "3-sigma" rule is conservative and designed to minimize risk from false alarms. In this application, the three levels on the gauges allow for an intermediate control limit at $k = 2$, corresponding to the amber region on the gauge. Addition of the intermediate control limit provides more rapid detection of faults [18]. Figure 2-9 shows how SPC maps the normal distribution to the green, amber and red regions of the gauges.

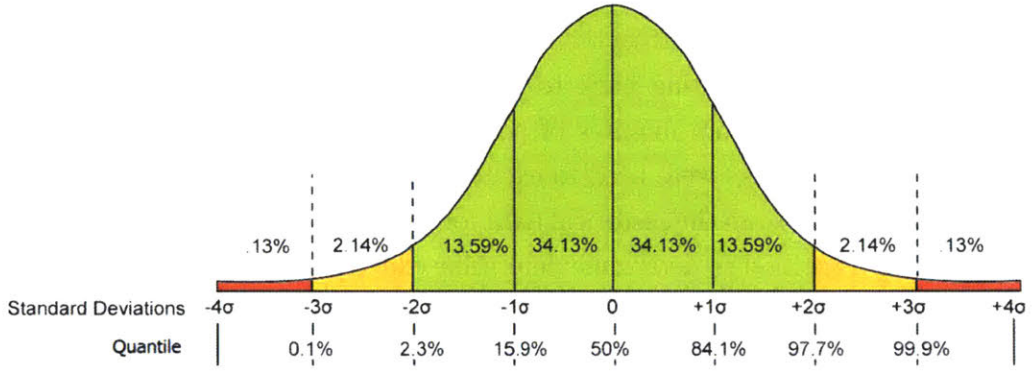

Figure 2-9: Progressive thresholds for fault detections [18]. Colors correspond to the red, amber, and green regions on the Dashboard gauges.

The SPC process can be adapted if the normal distribution does not properly fit the data. For example, the "Runs Per Day" metric for a controllable pitch propeller (CPP) pump is more accurately modeled by the Weibull distribution, often used in machinery reliability applications [17]. The probability density function for a two-parameter Weibull function is:

$$f(x) = (\frac{\beta}{\alpha})(\frac{x}{\alpha})^{\beta-1}(e^{\frac{x}{\alpha}})^{\beta} \tag{2.3}$$

where α is the scale factor and β is the shape factor. The shape factor and scale factor can be found by using the MATLAB *fitdist* function on historical "Runs Per Day" data (see Figure 2-10). To create the gauge regions for a non-normal distribution, the probability quantiles should match the red, amber, and green regions of the normal distribution in Fig. 2-9. The quantile (inverse cumulative distribution) function for a Weibull distribution is:

$$X(p, \alpha, \beta) = \alpha(-ln(1-q))^{\frac{1}{\beta}} \tag{2.4}$$

where q is the quantile. Therefore, the centerline can be found by setting q equal to .50 and solving for X. The upper and lower amber threshold levels can be found by setting q to .977 and .23, respectively. This ensures 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.

The statistical process control method is visualized below in Figure 2-10 for the WMEC-270 CPP pumps. For each of the five parameters, a histogram is created using the historical NILM data. The histogram is then modeled with multiple PDFs and the best fit function is selected. Depending on the PDF selected, the fault detection thresholds are set as shown in Figure 2-9 or using a quantile function as described in Equation 2.4. Chapter 4 provides case studies of fault conditions discovered using the parameters and fault detection thresholds described in this section.

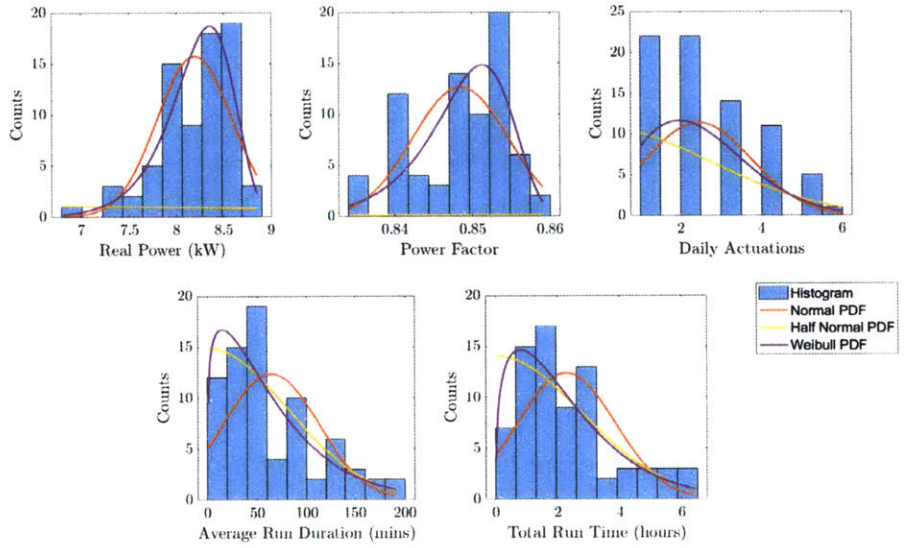

Figure 2-10: Different probability density functions fit to each parameter of the controllable pitch propeller (CPP) pumps.

2.4 Mission-Critical Equipment

This section provides a physical description of each piece of equipment that is monitored by NILM on the WMEC-270s. Specifically, this section describes the baseline electrical signatures used for load identification as well as the threshold levels used for fault diagnostics. Unless otherwise noted, the threshold levels are for SPENCER starboard panel equipment. Each piece of equipment should be analyzed individually to determine proper thresholds. This section expands upon information found in LT Gregory Bredariol's masters thesis [1].

2.4.1 MPDE Pre-Lube Pump

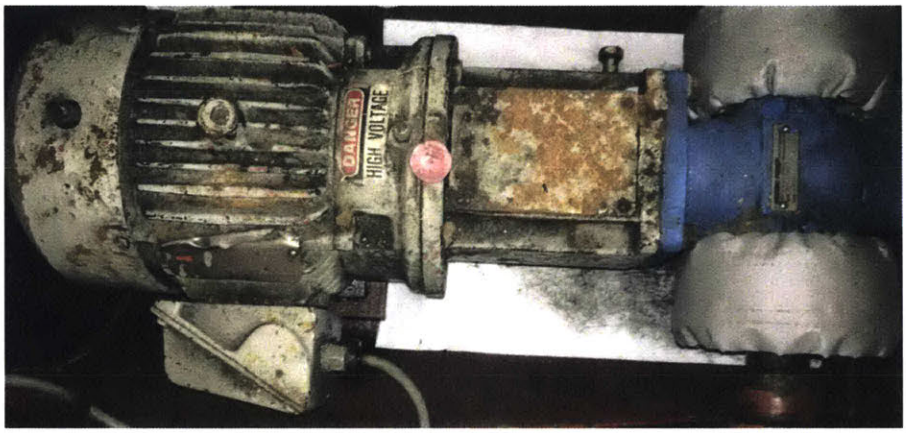

Figure 2-11: Main propulsion diesel engine pre-lube pump located in the engine room bilge on USCGC ESCANABA

The main propulsion diesel engine pre-lube pump is a motor-driven centrifugal pump that operates upon shutdown of the main diesel engines. The pump energizes once engine speed decreases below 150 RPM and secures when engine speed reaches 150 RPM. Above 150 RPM an attached gear-driven pump provides enough flow to effectively lubricate the engine [19,20]. Continuous circulation of lube oil prevents hot spots from developing and ensures that the engine is ready to start immediately [21]. Figure 2-12 shows the electrical signature of the pre-lube pump. Pre-lube pump normal behavior reflects that of the MPDE system and can often run for long periods of time when the engine is being kept in stand-by status. However, repeated short runs can indicate a control or sensing anomaly.

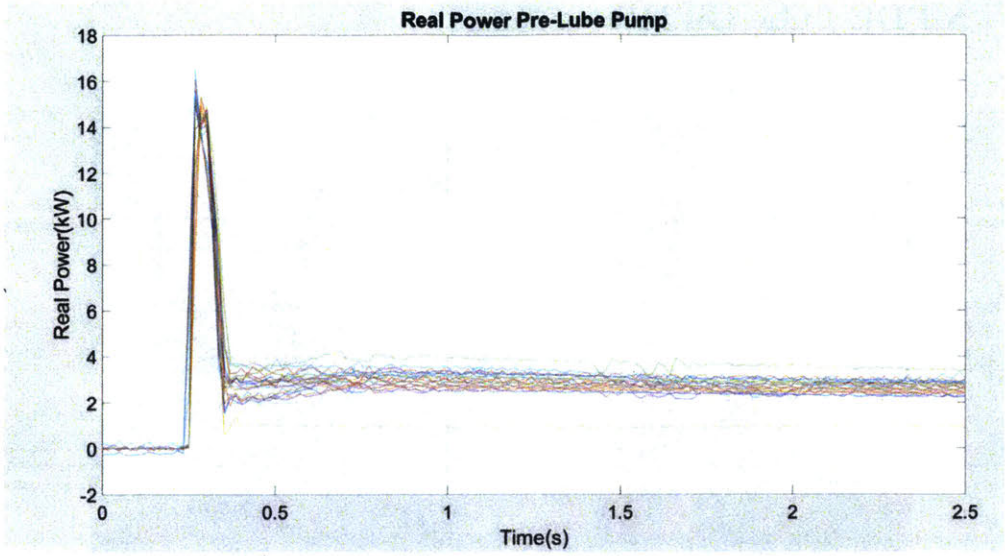

Figure 2-12: MPDE Pre-Lube Pump ON Transients

Panel(s):	Port and Starboard
Power Type:	3 - Phase
Steady-State Real Power (kW)	0.8
Steady-State Reactive Power (kVAR)	.4
Control Type	Automatic - Engine RPM

Table 2.1: MPDE Pre-Lube Pump Metrics

Power (kw)	Power Factor	Run Duration (mins)	Total Run Time (hrs)	Runs Per Day
0.7-1	.3-.5	2mins-8hrs	0 - 22	0 - 5
0.55-0.7/1-1.2	.25-.3/.5-.55	22-24	NA	5-6
-0.55/1.2+	-.25/.55+	10+	NA	6+

2.4.2 MPDE Lube Oil Heater

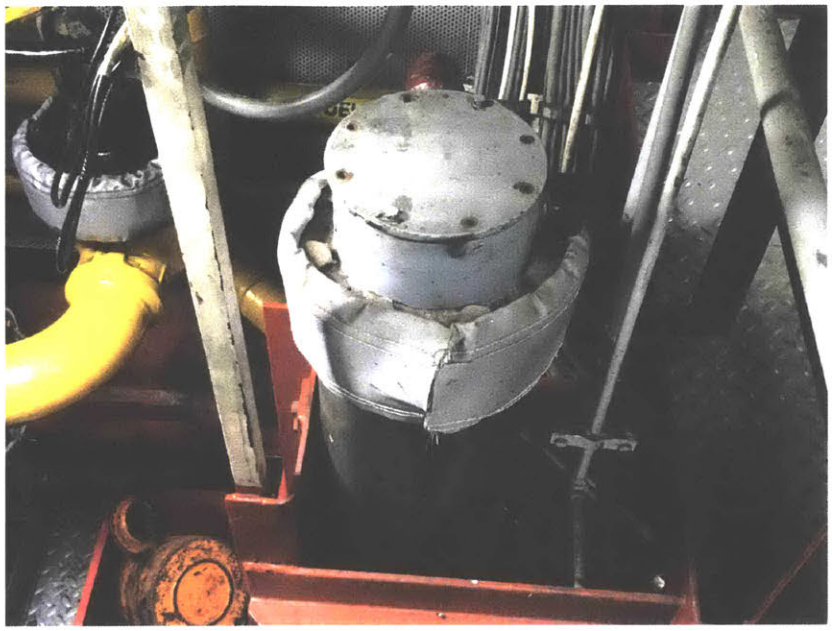

Figure 2-13: MPDE Lube Oil Heater on USCGC ESCANABA

The lube oil heater works in tandem with the pre-lube pump to maintain desired lube oil viscosity and temperature. The pre-lube pump described above circulates lube oil through the heater and then to the components of the engine. The purpose of this lube oil heater is to ensure that a minimum working temperature is maintained. Maintaining this temperature, especially during very cold operating conditions, can prevent excessive wear. If started cold, the lack of viscous lubricating fluid can cause metal-on-metal wear when the parts expand [21]. The heater secures upon engine start, then turns to automatic mode upon the engine securing [22]. In automatic mode the heater is thermostatically controlled to energize at the low temperature set-point (90F is standard) and secure at a high set-point (120F is standard) [20, 23]. The heater will not immediately energize upon engine stop due to entrapped heat and it may take several hours before enough heat is dissipated to require the heater to energize. The heater can run for hours at a time, especially in a cold engine room. The heater shares a thermostatic control switch with the MPDE jacket water heater [22]. When the engine is in stand-by status, these loads operate together, appearing as a 21kW, purely resistive load (Fig. 2-14). Runs of less than 5 minutes or runs of more than 6 hours can indicate a control fault is present.

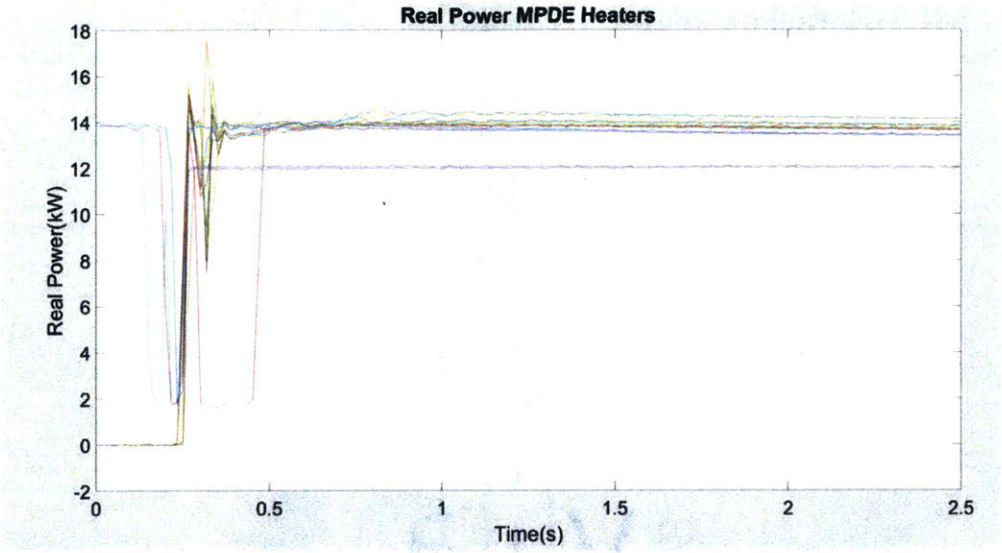

Figure 2-14: Transients from the MPDE lube oil and jacket water heaters. These two loads often actuate in tandem. In the majority of these transients, the jacket water heater is operating in reduced capacity, due to corrosion detailed in Chapter 4.

Panel(s):	Port and Starboard
Power Type:	3 - Phase
Steady-State Real Power (kW)	12
Steady-State Reactive Power (kVAR)	0
Control Type	Automatic - Thermostat

Table 2.2: MPDE Lube Oil Heater Metrics

Power (kw)	Power Factor	Run Duration (mins)	Total Run Time (hrs)	Runs Per Day
10.1-11.7	.9995-1	2mins-8.5hrs	0 - 20	0 - 6
9-10.1/11.7-12.5	NA	1-2mins/8.5-10hrs	20-24	6-7
-9/12.5+	-.9995/1+	.5-1/10+	NA	7+

2.4.3 MPDE Jacket Water Heater

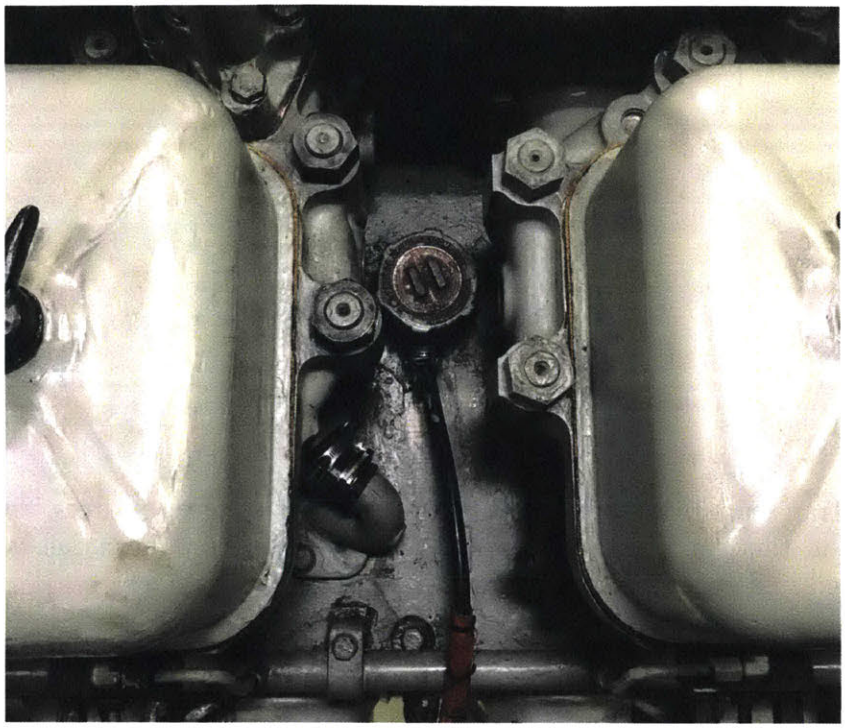

Figure 2-15: MPDE jacket water heater on USCGC ESCANABA

The final component used to maintain MPDE readiness is the jacket water heater. This load consists of two 4.5kW three-phase immersion heating elements sitting between the cylinders on either side of the engine block. These elements are served from the same breaker and controller, and appear to NILM as a single 9kW resistive load. The heaters are thermostatically controlled to energize at a low set point (90F is standard), and secure at a high set-point (120F is standard) [22]. Jacket water is circulated by a separate 120V electric pump (not monitored) [23]. Jacket water is distilled water with additives included to prevent corrosion and freezing. It is designed as an intermediate heat transfer medium, separating lube oil from cooling seawater to ensure that a leak of lube oil to the sea is highly unlikely. Lube oil cools the engine directly, then jacket water transfers heat from the lube oil to seawater [21]. The jacket water heater ensures that the engine does not fully cool to "cold iron" status.

Similar to the lube oil sump heater described above, this heater remains secured for hours after engine shutdown due to latent heat. When underway, the heater operates in tandem with the lube oil heater, as shown in Fig. 2-14. During extended inports the lube oil heater is secured and the jacket water heater is kept running to ensure the engine does not

reach "cold iron" status. Despite the non-corrosive nature of jacket water, NILM detected a significant corrosion problem with these heaters, detailed in depth in Chapter 4. This fault caused significant variation in the electric signature of the heater, so a plot of representative transients is not included here. Due to the ongoing fault condition, the fault thresholds below are based on the limited healthy operations observed.

Panel(s):	Port and Starboard
Power Type:	3 - Phase
Steady-State Real Power (kW)	9
Steady-State Reactive Power (kVAR)	0
Control Type	Automatic - Thermostat

Table 2.3: MPDE Jacketwater Heater Metrics

Power (kw)	Power Factor	Run Duration (hrs)	Total Run Time (hrs)	Runs Per Day
8.3-9.4	.995-1	15mins-5hrs	0 - 16	0 - 4
8.0-8.3/9.4-9.7	.99-.995	5-11	16-22	4-8
-8.0/9.7+	-.99/1+	11	22+	8+

2.4.4 SSDG Lube Oil Heater

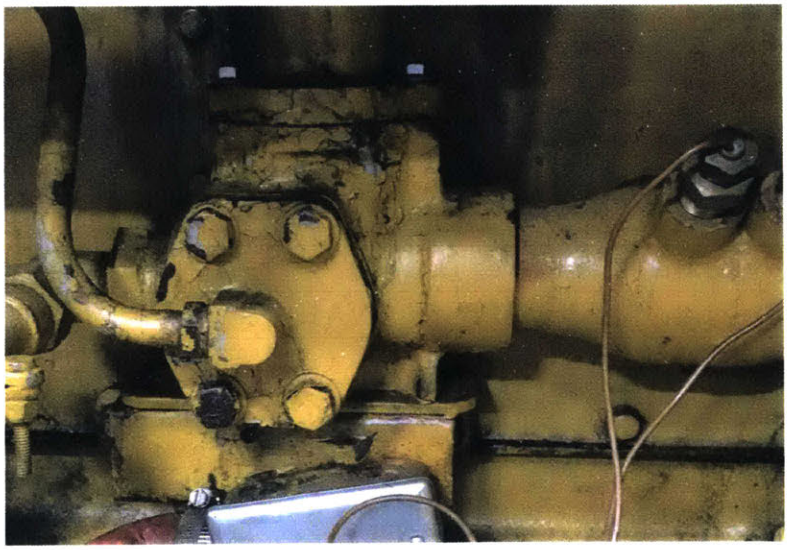

Figure 2-16: Ship's Service Diesel Engine (SSDG) Lube Oil heater on USCGC ESCANABA

Similar to the MPDE lube oil heater, the SSDG lube oil heater keeps engine lubricating oil temperature within a set range while the engine is secured [22]. For this system, the prescribed range is a low of 90F and a high of 120F. Unlike the other 3-phase loads served from the sub-panels, the lube oil heater is a line-to-line single-phase load of 1.3kW. A typical SSDG lube oil heater transient is shown below (Figure 2-17). A change is real power is detected on only two phases, as well as off-setting positive and negative changes in reactive power.

The heater quickly secures when the engine is brought online and the temperature rapidly increases above the upper set point. The lube oil heater operates in relatively short cycles of 10-30 minutes when the engine is secured.

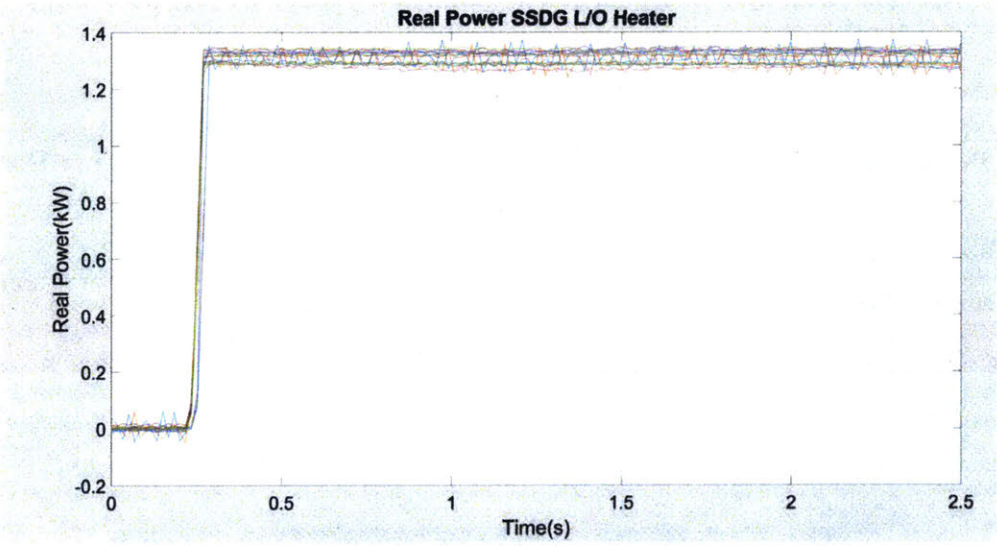

Figure 2-17: SSDG Lube Oil Heater ON Transients

Panel(s):	Port and Starboard
Power Type:	2 - Phase Line-to-Line
Steady-State Real Power (kW)	1.3
Steady-State Reactive Power (kVAR)	± .9
Control Type	Automatic - Thermostat

Table 2.4: SSDG Lube Oil Heater Metrics

Power (kw)	Power Factor	Run Duration (mins)	Total Run Time (hrs)	Runs Per Day
1.24-1.32	.998-1.005	1-6	0 - 19	0 - 214
1.2-1.24/1.32-1.35	NA	.5-1/6-7	19-23	214-277
-1.2/+1.35	-.998/1.005+	-.5/+7	23+	277+

2.4.5 SSDG Jacket Water Heater

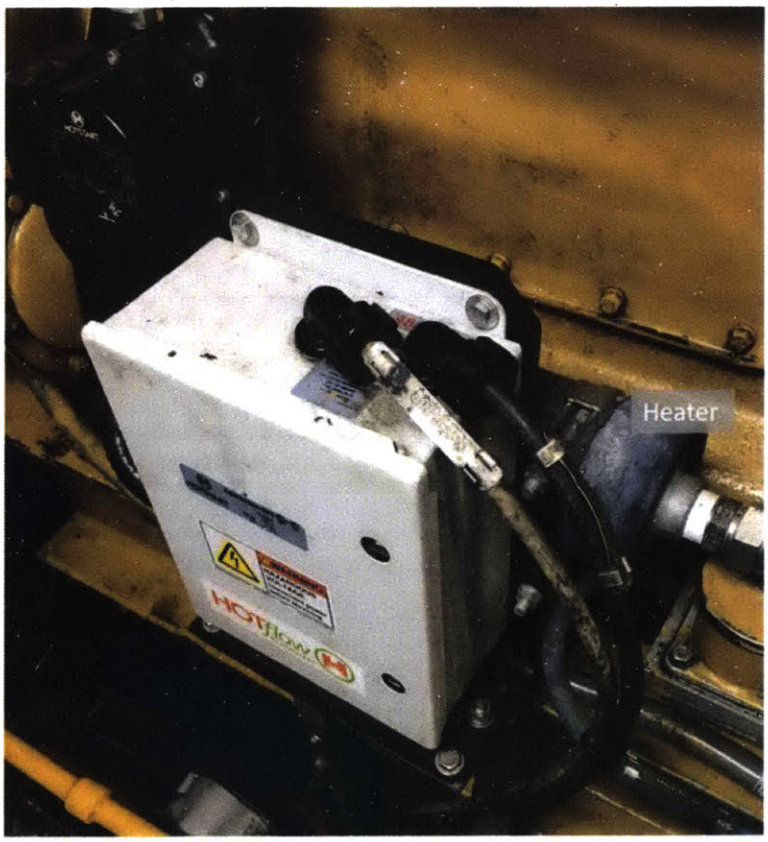

Figure 2-18: Ship's service diesel generator jacket water system on USCGC ESCANABA, including in-line pump and heater. Only the heater is monitored.

The SSDG jacket water heater serves the same purpose as the MPDE jacket water heater, maintaining the water temperature in a set range when the engine is secured. The thermostatic range for this element is 85F to 115F [22]. The heater is in-line to fluid flow and works in tandem with a circulation pump (not monitored). Its electric signature and behavioral metrics are shown below.

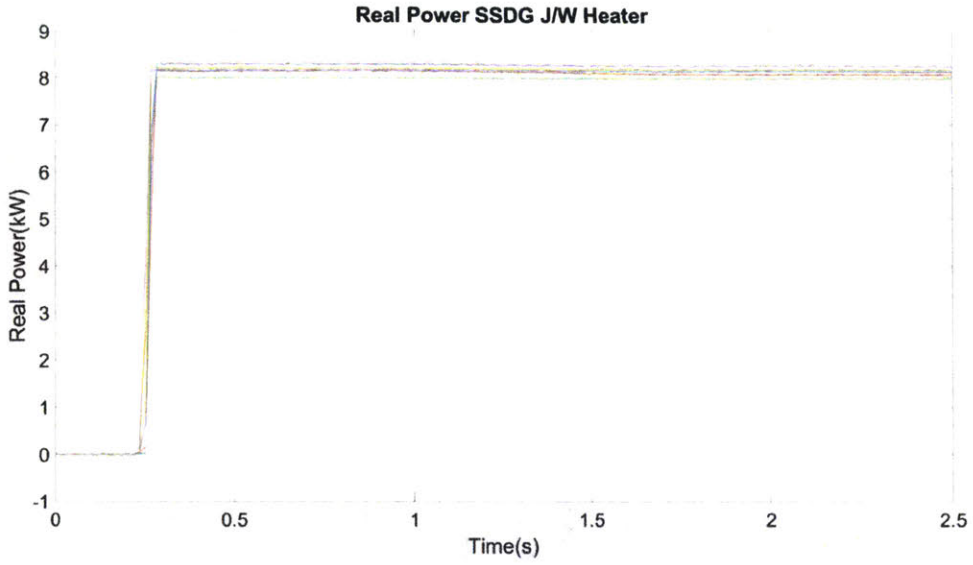

Figure 2-19: SSDG Jacketwater Heater ON Transients

Panel(s):	Port and Starboard
Power Type:	3-Phase
Steady-State Real Power (kW)	7.5
Steady-State Reactive Power (kVAR)	0
Control Type	Automatic - Thermostat

Table 2.5: SSDG Jacketwater heater metrics

Power (kw)	Power Factor	Run Duration (mins)	Total Run Time (hrs)	Runs Per Day
7.8-8.4	.999-1	13mins-3hrs	0 - 15	0 - 7
7.6-7/8.4-8.5	NA	1-13mins/3-4hrs	15-18	7-10
-7.6/8.5+	-.99/1.001+	-1min/4hrs+	18+	10+

2.4.6 Graywater Pumps

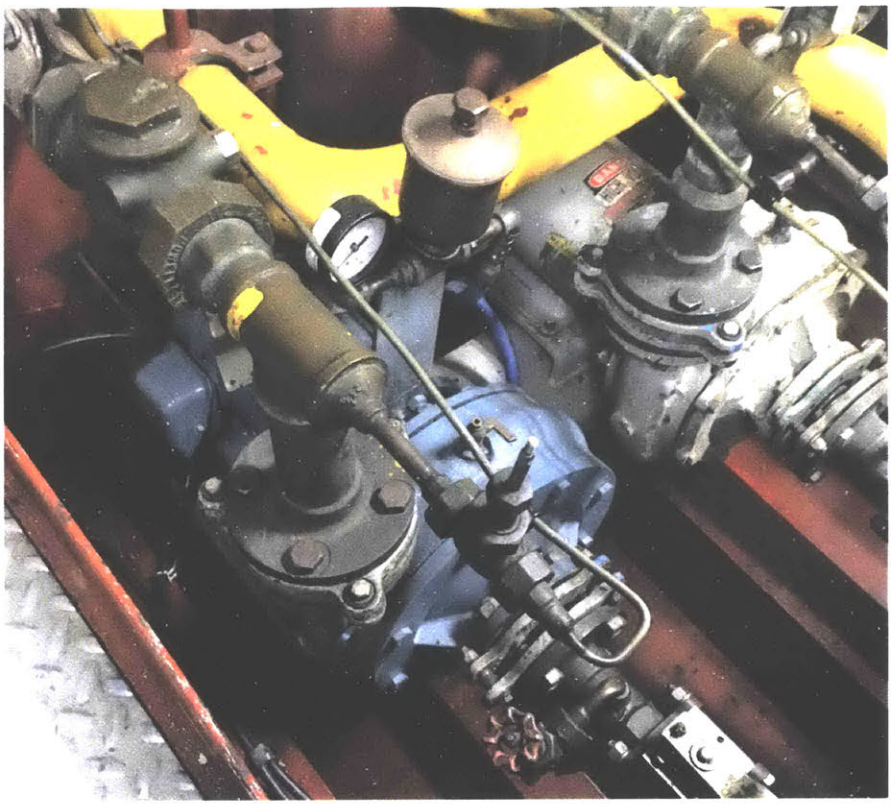

Figure 2-20: Graywater system pumps on USCGC ESCANABA.

The graywater system in the main engine room consists of a gravity-fed 138-gallon (522-liter) collection tank and two transfer pumps [24]. The tank fills from showers and sinks around the aft portion of the vessel. Two identical motor-driven centrifugal pumps alternate runs to empty the tank when a high level is reached. Conductivity sensors detect water levels and provide feedback for pump control. The high sensor set point is the 92-gallon mark and the low level sensor set point is the 13-gallon mark [24]. Figure 2-21 shows the key elements of this system. At the pier, the wastewater is pumped to a shore-side connection. When steaming offshore, the wastewater is pumped directly over the side. In waters close to shore or in protected marine areas, the wastewater is pumped to a larger holding tank. A representative electronic signature for these pumps is shown in Fig. 2-22 below.

From the sensor set points, a pump discharges 79 gallons of gray water each cycle. From estimates of the head pressure in the system and manufacturer-provided pressure-flow rate curves, the flow rate should be approximately 60 gallons per minute [24]. This yields a typical expected pump run duration of about 80 seconds, which is confirmed from observations of

NILM data. The tank fills at very different rates depending on the status of ship's crew. During busy periods underway the pumps will run multiple times per hour, but during a slow day inport the pumps will only run a few times per day. The tank-level sensors that control pump operations are susceptible to failure, leading to either very short or extremely long pump runs. NILM has detected multiple faults in this system, detailed in Chapter 4,

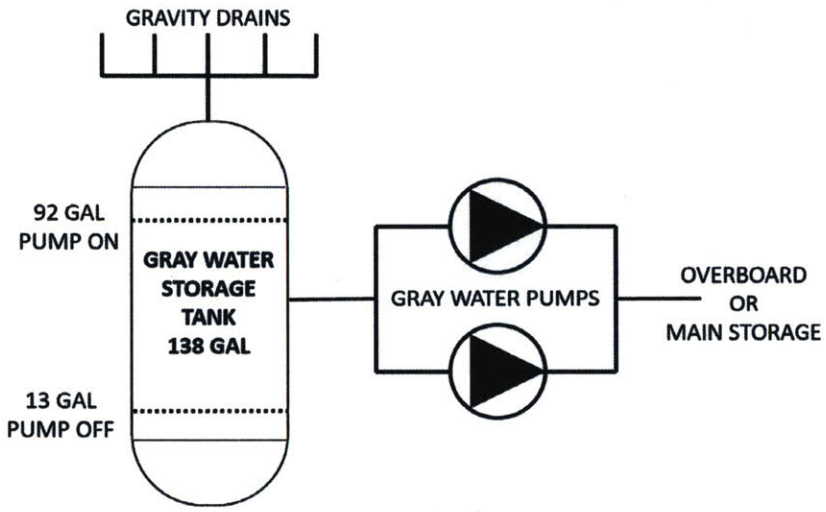

Figure 2-21: Graywater holding and transfer system on WMEC-270 cutters. Pumps automatically alternate operation on each cycle.

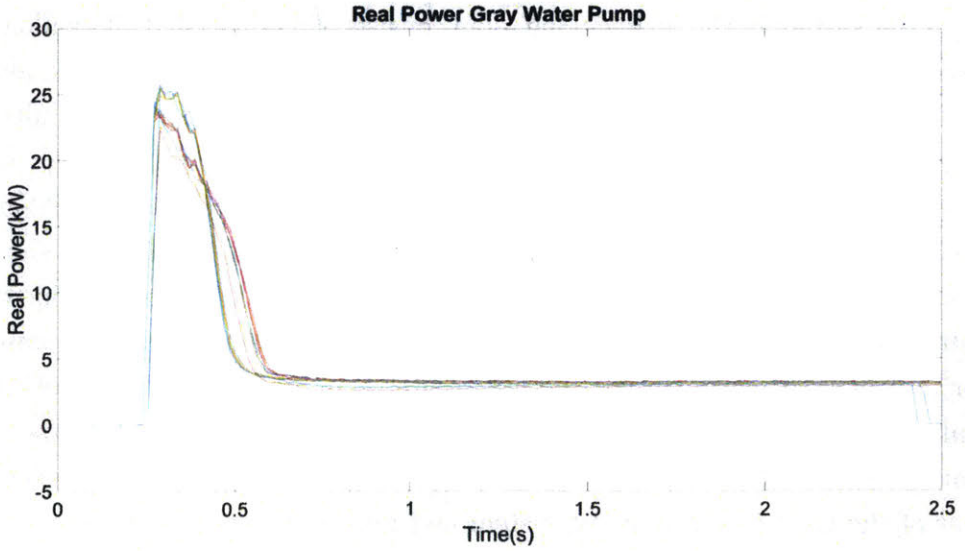

Figure 2-22: Graywater ON transients

Panel(s):	Port
Power Type:	3-Phase
Steady-State Real Power (kW)	2.9
Steady-State Reactive Power (kVAR)	2.4
Control Type	Automatic - Tank Level

Table 2.6: Graywater Pump Metrics

Power (kw)	Power Factor	Run Duration (mins)	Total Run Time (hrs)	Runs Per Day
2.65-2.88	0.7 - 0.76	.88 - 1.2	0.05 - .5 hrs	3 - 29
2.6-2.65/2.88-2.89	.85-.88/1.2-1.22	.03-.05/.5-.57	1-2.5	2-3/29-31
-2.6/2.89+	-0.69/.77+	-0.85/1.22+	.03-/.57+	-2/31+

2.4.7 Controllable Pitch Propeller Hydraulic Pump

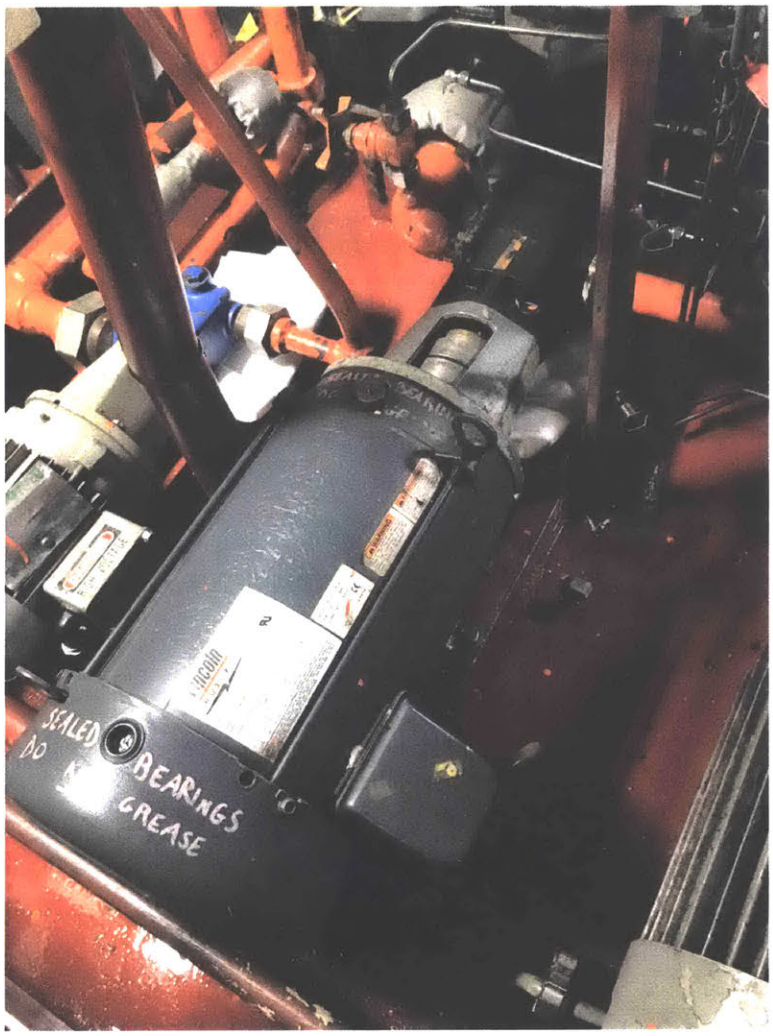

Figure 2-23: Controllable Pitch Propeller Hydraulic Pump

Controllable pitch propellers (CPPs) allow users to adjust the amount of thrust generated by a propulsor while maintaining a fixed rotational speed. The blades of a propeller are moved via a hydraulic control system that runs from the engine room to the propellers through the center of the propulsion shaft [20]. In addition to increasing the maneuverability of the ship, the CPP system allows for high-propeller-speed, low-thrust propulsion profiles that are essential to eliminate under-loading of diesel engines when traveling at low speed [21]. The WMEC-270 CPP system uses three pumps to maintain hydraulic control pressure at the propeller hubs. Under normal operating conditions, pressure is supplied by the main hydraulic pumps (CPP pumps "A" and "B"); however, if flow drops below a set point (44

gpm), the monitored pump (pump "C") is energized [25].

The "C" pump is also energized when the restricted maneuvering doctrine (RMD) has been set, signaling that the ship has entered a heightened state of readiness [20]. These heightened states include transiting through restricted waters, entering and leaving port, a general emergency, law enforcement operations, flight operations, small boat operations, or any other time the commanding officer deems an increased readiness posture is required. During these operations, a failure of the system can be catastrophic; therefore, these "C" pumps are energized to improve redundancy in the system and manage risk. Because these pumps are operated either manually or in emergencies, their behavior has a large variance. The normal electrical signature and behavior parameters are shown below.

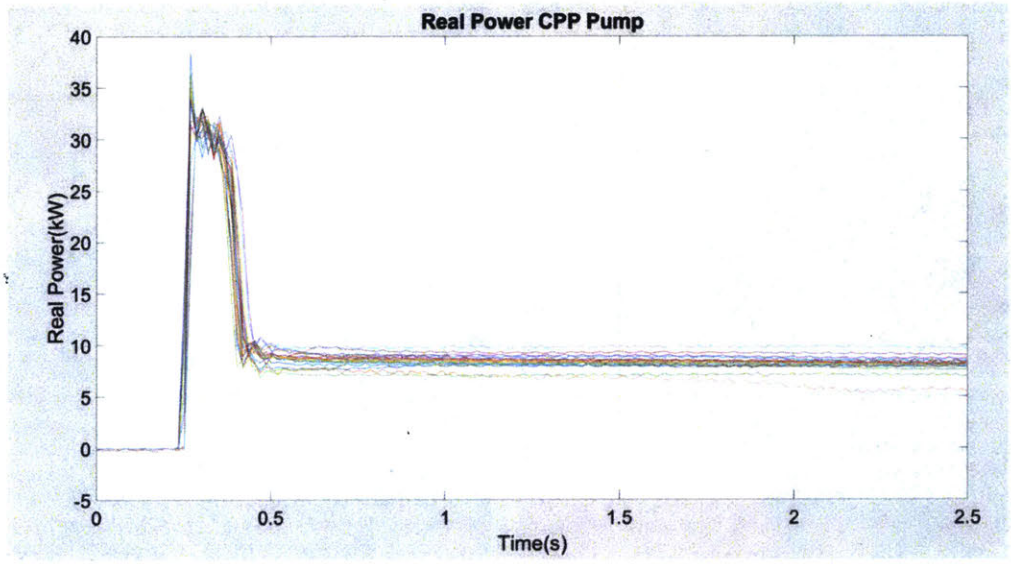

Figure 2-24: Controllable Pitch Propellor ON transients

Panel(s):	Port and Starboard
Power Type:	3-Phase
Steady-State Real Power (kW)	6.2
Steady-State Reactive Power (kVAR)	4.8
Control Type	Manual or Automatic (Pressure)

Table 2.7: CPP "C" Pump Metrics

Power (kw)	Power Factor	Run Duration (mins)	Total Run Time (hrs)	Runs Per Day
5.5-8	0.78 - 0.85	7-170	0 - 5.5	0 - 5
5-5.5/8-8.5	.76-.78/.85-.87	1-3/170-190	5-6	5-6
-5/8.5+	-.76/.87+	-1/190+	6+	6+

2.4.8 Fuel Oil Purifier

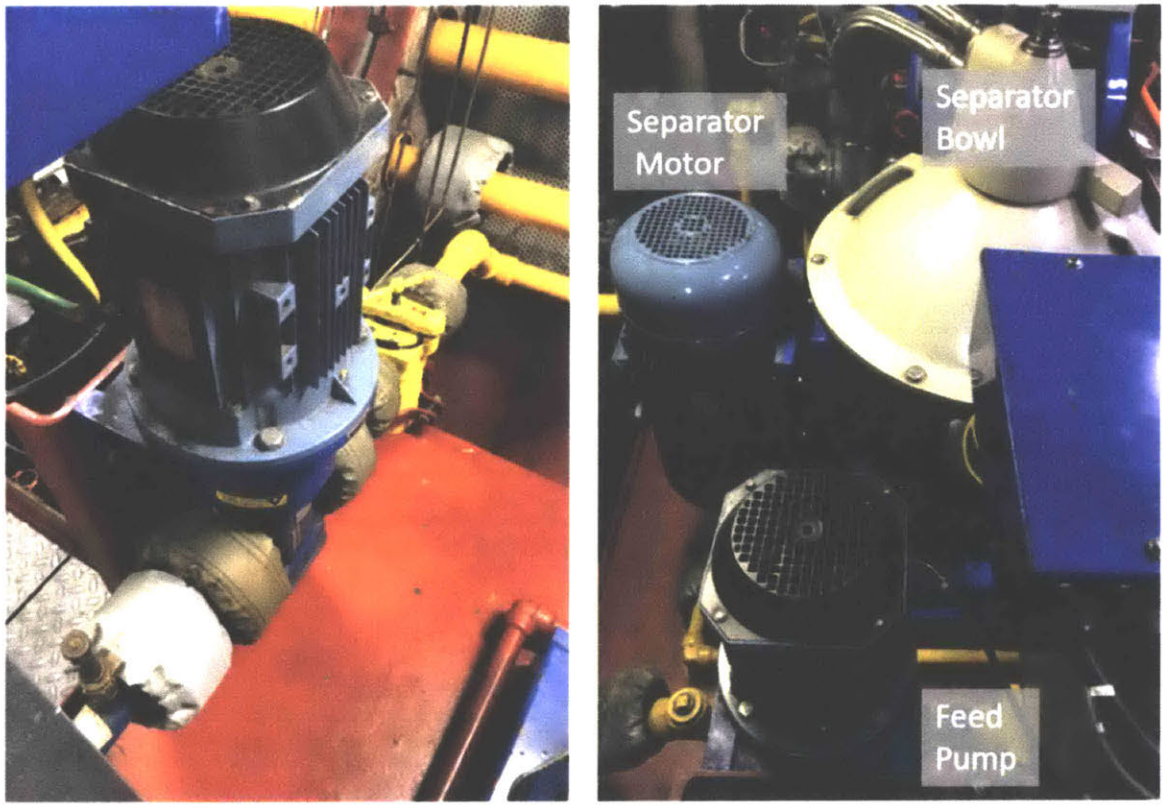

Figure 2-25: Left: Fuel Oil Purifier feed pump. Right: FOP system including feed pump, separator motor, and separator bowl.

The fuel oil purifier (FOP) cleans diesel oil before use in the MPDE or SSDG. The purifier system draws fuel from various storage tanks around the ship via a feed pump, and processes the fuel using a motor-driven 9000 RPM centrifugal separator system [26]. Both elements are fed from the same breaker on the port panel and can be detected by NILM. The system also includes a heater that is fed from a separate breaker on the main switchboard and is not

monitored. The purifier is run frequently when the ship is underway and increases in use if poor quality fuel is onboard. The process of operation goes through multiple steps. These steps are managed by a programmable logic controller (PLC), common for more complex shipboard systems. Once the valves are aligned, the operator simply presses the start button and the process begins. NILM sensors detect each of these steps, depicted below in Figure 2-26.

Like the controllable pitch propeller pumps, the FOP is controlled manually, leading to a wider variance in system behavior. The operational metrics for both the feed pump and centrifugal separator are shown below.

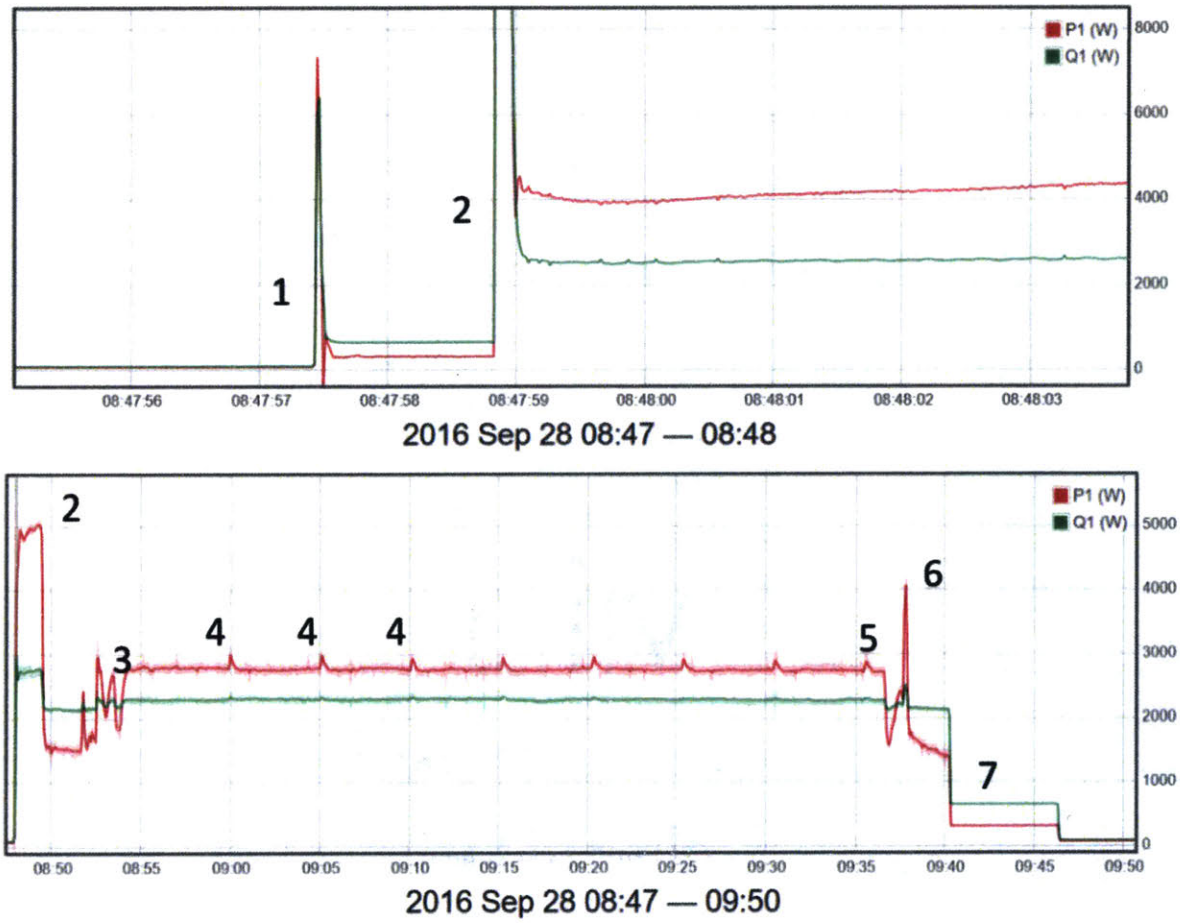

Figure 2-26: Power data showing fuel oil purifier operating sequence. Top: Start sequence. Bottom: Full Operation Cycle of FOP

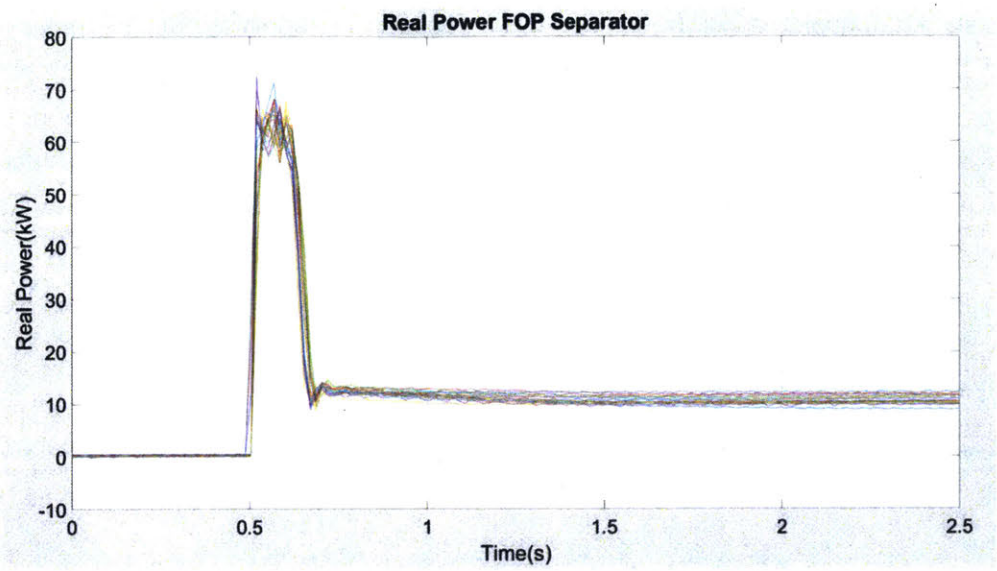

Figure 2-27: Fuel Oil Purifier separator ON transients.

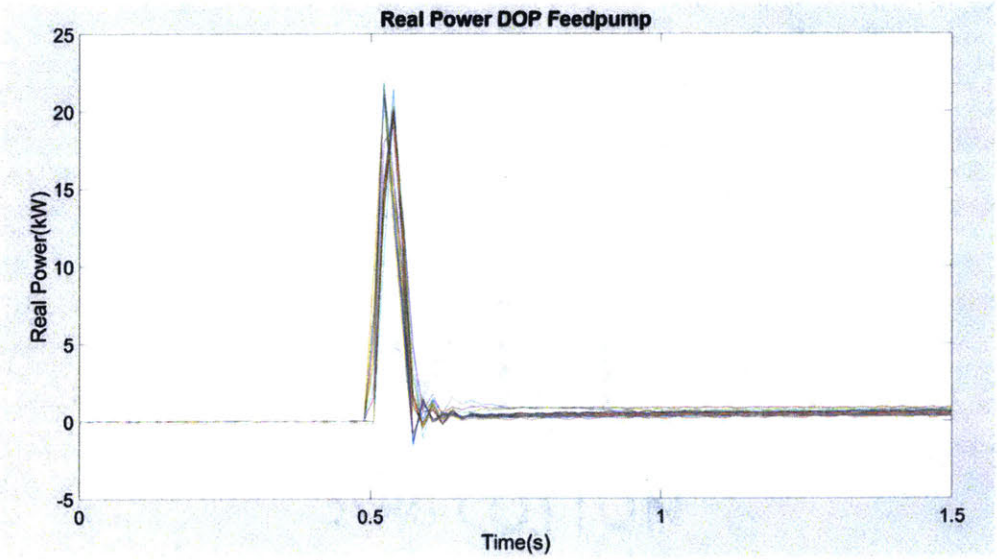

Figure 2-28: Fuel Oil Purifier feed pump ON transients.

Panel(s):	Port
Power Type:	3-Phase
Steady-State Real Power (kW)	Feedpump - .95 Centrifuge - 9
Steady-State Reactive Power (kVAR)	Feedpump - 1.9 Centriguge - 6.3
Control Type	Manual

Table 2.8: FOP Separator Metrics

Power (kw)	Power Factor	Run Duration (mins)	Total Run Time (hrs)	Runs Per Day
1.9-14	.57-.93	2-32	0-.9	0-5
1-1.9/14-17	.48-.57/.93-1	1-2/32-41	.9-1.1	5-7
-1/+17	-.48/+.1	-1/41+	1.1+	7+

Table 2.9: FOP FeedPump Metrics

Power (kw)	Power Factor	Run Duration (mins)	Total Run Time (hrs)	Runs Per Day
.59-.82	.36-45	24-50	0 - 1	0 - 5
.56-.59/.82-.84	.34-.36/.45-.46	21-24/50-51	1-1.5	5-7
-.56/.84+	-.34/.46+	-21/51+	1.5+	7+

2.4.9 Bilge and Emergency Ballast Pump

Figure 2-29: Bilge and emergency ballast pump on USCGC SPENCER

The bilge and emergency ballast (BEB) pump is motor-driven centrifugal pump used for emptying the engine room bilges of excess water and also for taking on ballast water for stability purposes. When the ship has burned a significant portion of its fuel load, ballast water is taken on to maintain safe stability and a smooth ride for the crew. Despite its name, the pump is not used only in emergencies. Bilges are pumped at least once per week, and ballast water is frequently added and discharged when underway. Water for ballast is drawn in via the seachest and can be pumped to any of six dedicated ballast tanks [27]. This system can also be aligned to remove water from each tank and empty over the side [27]. Typical operation of the BEB pump is characterized by multiple successive runs. The pump must be secured when switching suction/fill from one tank to the next. A typical pump transient and operational metrics for the pump are shown below.

The bilge and emergency ballast pump's electrical signature presents some unique challenges for characterization and classification. Unlike most loads on the panels, the time it takes the pump to reach the steady state is highly variable. This variability can likely be attributed to air pockets within the bilge and ballast pumping system. When pumping out bilges and ballast tanks, operators will try to get the tanks and bilges to the lowest level possible, causing the pump to take in a mixture of air and water. When the pump is then secured and suction is shifted to a new tank, this air remains in the system and leads to a prolonged start sequence. Figure 2-30 shows two runs of the BEB pump, highlighting the differences in possible ON transients.

Despite being operated manually, BEB pump run frequency and duration is consistent; tanks and bilges should roughly take the same amount of time to empty and fill. The behavioral metrics for the BEB pump are shown below.

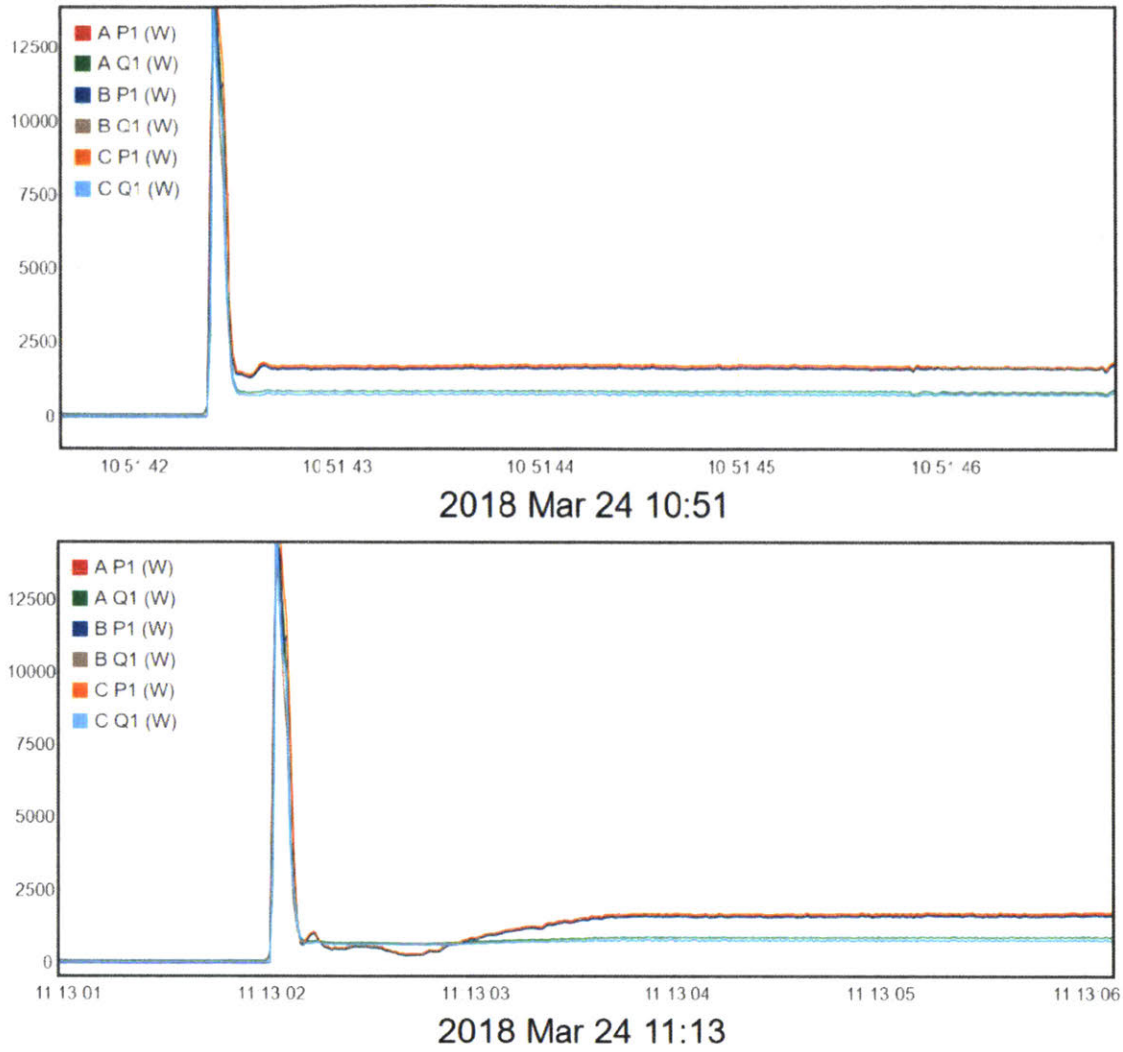

Figure 2-30: Top: Bilge and emergency ballast pump normal start sequence. Bottom: BEB Pump displaying extended start sequence.

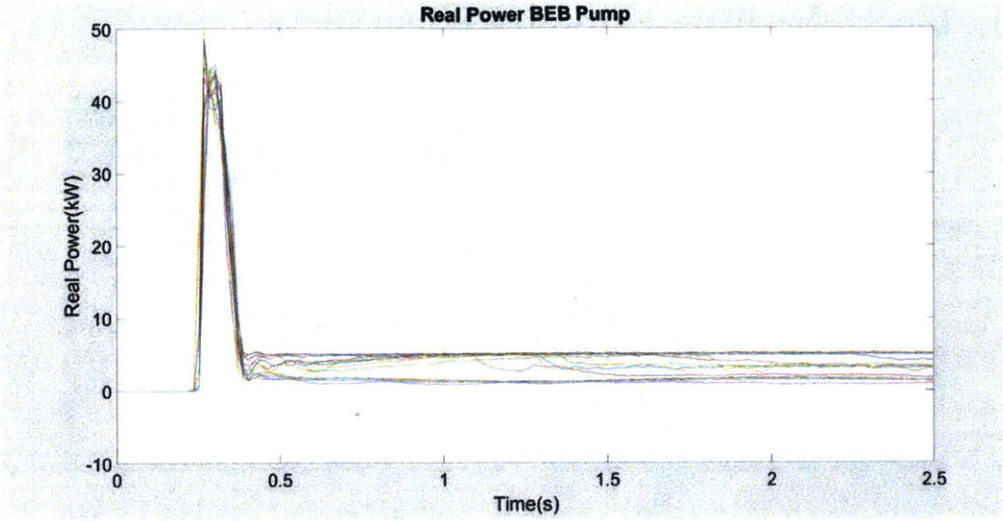

Figure 2-31: Bilge and emergency ballast pump ON transients

Panel(s):	Starboard
Power Type:	3-Phase
Steady-State Real Power (kW)	5.1
Steady-State Reactive Power (kVAR)	2.6
Control Type	Manual

Table 2.10: Bilge and Emergency Ballast Pump Metrics

Power (kw)	Power Factor	Run Duration (mins)	Total Run Time (hrs)	Runs Per Day
1.5-6	.6-1	3-30	0-2	0 - 7
.2-1.5	.5-.6/1-1.1	2-3/33-35	2-3	7-8
-.2/+7.5	-.5/1.1+	-2/35+	3+	8+

2.4.10 Inport Auxiliary Saltwater Pump

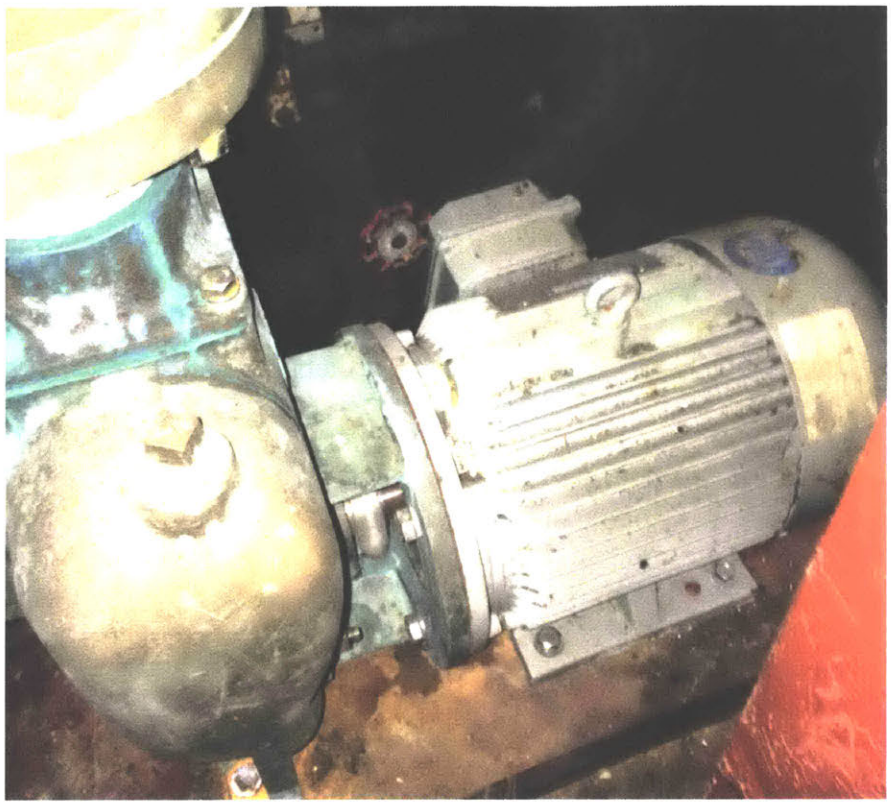

Figure 2-32: Inport ASW pump on USCGC SPENCER

The inport auxiliary saltwater (ASW) pump provides cooling saltwater when the cutter is at the pier. When at sea, the cutter uses two 30kW pumps (not monitored) to deliver salt water to engines, generators and a variety of other critical systems [20]. In port the only loads that require cooling water are the air conditioning and refrigeration units. This pump operates continuously for the days and weeks when the cutter is at the pier. Due to the limited observed operations, accurate fault thresholds cannot be determined at this time. Any repeated cycling of this load is evidence of a fault condition. The normal ON transients for the pump are shown below.

Panel(s):	Port
Power Type:	3 - Phase
Steady-State Real Power (kW)	3.5
Steady-State Reactive Power (kVAR)	2
Control Type	Manual

2.4.11 Shaft Turning Gear Motor

Figure 2-33: Shaft turning gear motor on USCGC SENECA.

The turning gear motor is used to slowly rotate the propulsion shafts to bring propellers and shafts to a desired position for maintenance, usually when the ship is in drydock. The gear is clutched into the shaft and manually jogged to precisely rotate the shaft [20]. The operation of this motor is very rare and has not been observed by NILM.

Panel(s):	Stbd
Power Type:	3 - Phase
Steady-State Real Power (kW)	1.5 (estimated)
Steady-State Reactive Power (kVAR)	1 (estimated)
Control Type	Manual

2.4.12 Oily Water Separator

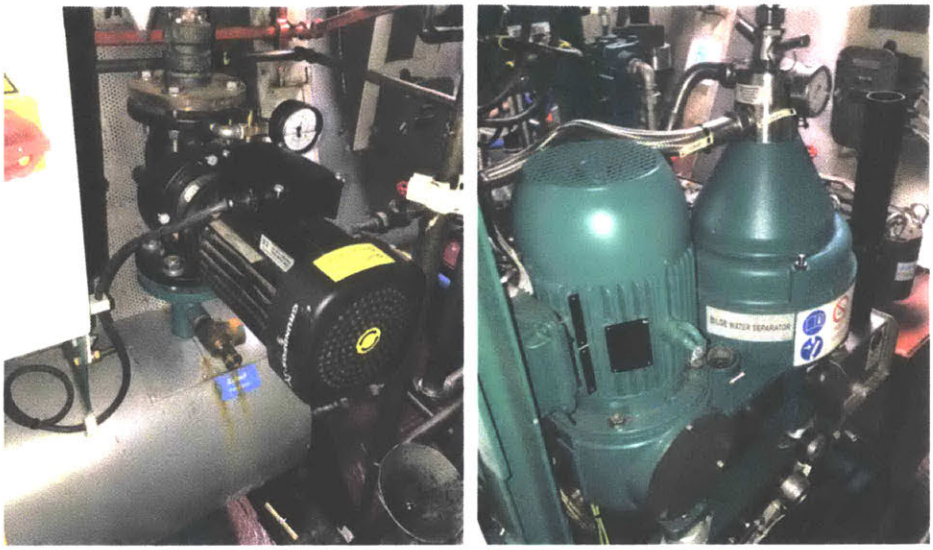

Figure 2-34: Left: Oily water separator feed pump. Right: OWS centrifugal separator.

The oily water separator (OWS) is newly-installed centrifugal purifying system that is similar in design and operation to the FOP. This system replaces a defunct gravity-coalescer system. The oily water separator processes oily waste and removes oil so clean water can discharged to the ocean. This system reduces the amount of waste water stored, allowing the cutter to remain at sea for longer periods. The system is regulated by an oil content meter that ensures water being discharged is within environmental regulations. The new OWS skid consists of a feed pump, motor-driven separator, and an additional filtration system (Figure 2-34). The OWS has only been run for testing and evaluation and has not been observed by NILM.

Panel(s):	Stbd
Power Type:	3 - Phase
Steady-State Real Power (kW)	TBD
Steady-State Reactive Power (kVAR)	TBD
Control Type	Manual

Chapter 3

Automatic Watchstanding

Non-intrusive load monitoring can reduce the burden on watchstanders by automating time-intensive machinery watchstanding and assisting with equipment fault detection. Lowering the workload of sailors has never been important than today. In recent years, the U.S. Navy experienced a series of deadly, high-profile navigation incidents. Subsequent investigations found a common problem on all the affected vessels: overworked and under-trained crew members [28]. Automation and sensing technology can be part of the solution. Increased automation at sea can minimize paperwork and bureaucracy, assist with predictive and preventative maintenance, provide decision-making support, and reduce or even eliminate some repetitive tasks [29].

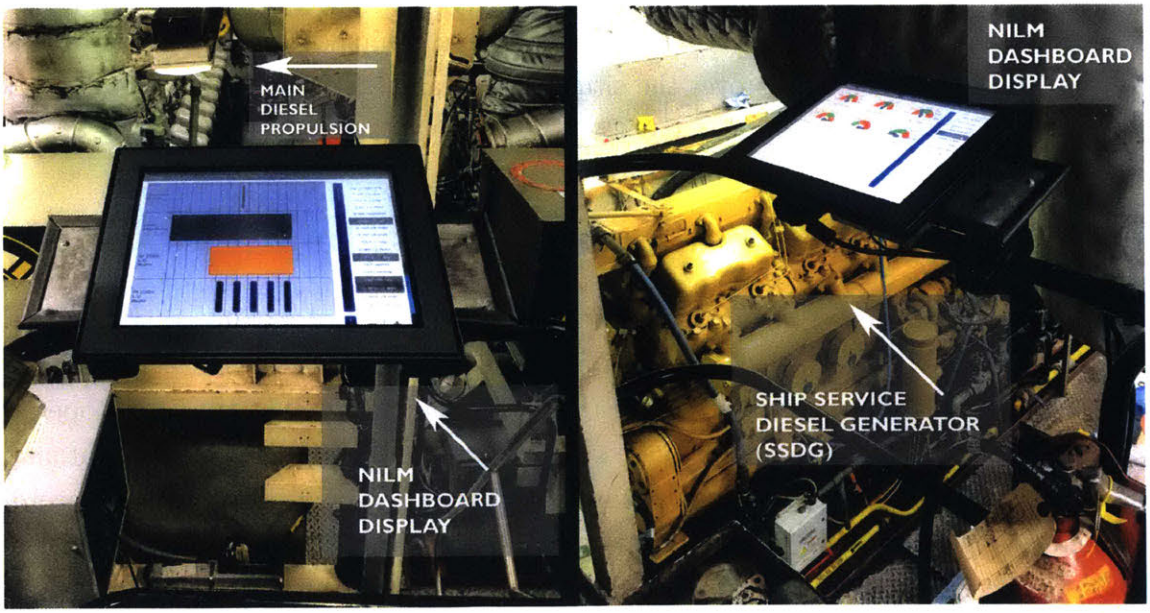

Figure 3-1: The Non-intrusive Load Monitoring Dashboard, shown here on USCGC SPENCER, can improve the situational awareness of watchstanders on the WMEC-270.

The movement towards increased automation is well underway in the machinery spaces of modern ships, where the proliferation of centralized machinery control and monitoring systems (MCMS) have revolutionized the way that sailors stand machinery watch. The ability to view the status of equipment across the ship from a single point is a powerful situational awareness tool. MCMS systems reduce manual logkeeping effort and allow for fewer watchstanders to be required at any given time. Most new ships are built from the ground up with a large sensor network to support centralized MCMS.

For vessels that are already in service, gaining the benefit of centralized monitoring often requires an expensive and lengthy installation of a large sensor network. Legacy equipment is often not designed to interface with modern sensing and control systems. However, the information needed to create a machinery monitoring network is already present in the ship's power system. By disaggregating equipment operating schedules from power data, described in the previous chapter, NILM can create an effective centralized monitoring point with a few well-placed sensors. Non-intrusive load monitoring can serve as an "automatic watchstander" for crews on legacy vessels, reducing repetitive manual logkeeping and improving crew situational awareness about equipment behavior.

The NILM Dashboard, shown above on SPENCER in Figure 3-1, is designed to act as an automatic watchstander for the crew of the WMEC-270. These ships are an excellent example of a legacy asset that can benefit from the addition of a centralized monitoring point, provided that the information is presented in a clear, intuitive manner. The Timeline view of the Dashboard, shown in Figure 1-4 and in the sections below, provides the crew with insight into the operational history of a variety of unmonitored equipment. This chapter describes how the NILM Dashboard improves the situational awareness of machinery watchstanders on the WMEC-270.

3.1 Determining Engine and Ship Status

The previous chapter describes the equipment served by the two monitored sub-panels on SPENCER and ESCANABA (see Table 1.1 for a list of equipment). By disaggregating individual load events from the overall power stream in each panel, NILM tracks the operation of 20 pieces of equipment. Many of these systems are operated by closed-loop, automated controllers such as temperature switches and tank level sensors. Tracking and recording the operation of all of these systems is a nearly impossible task using manual logging techniques, but their behavior is captured and clearly displayed by NILM. Figure 3-2 shows how NILM transforms power data into a log of machinery operations. From the NILM Dashboard "Timeline" view shown below, the ship's crew can view the present status and the previous

operating schedules of each load. The timeframe and loads of interest are selected by the user, with the colored bars representing periods when equipment is energized.

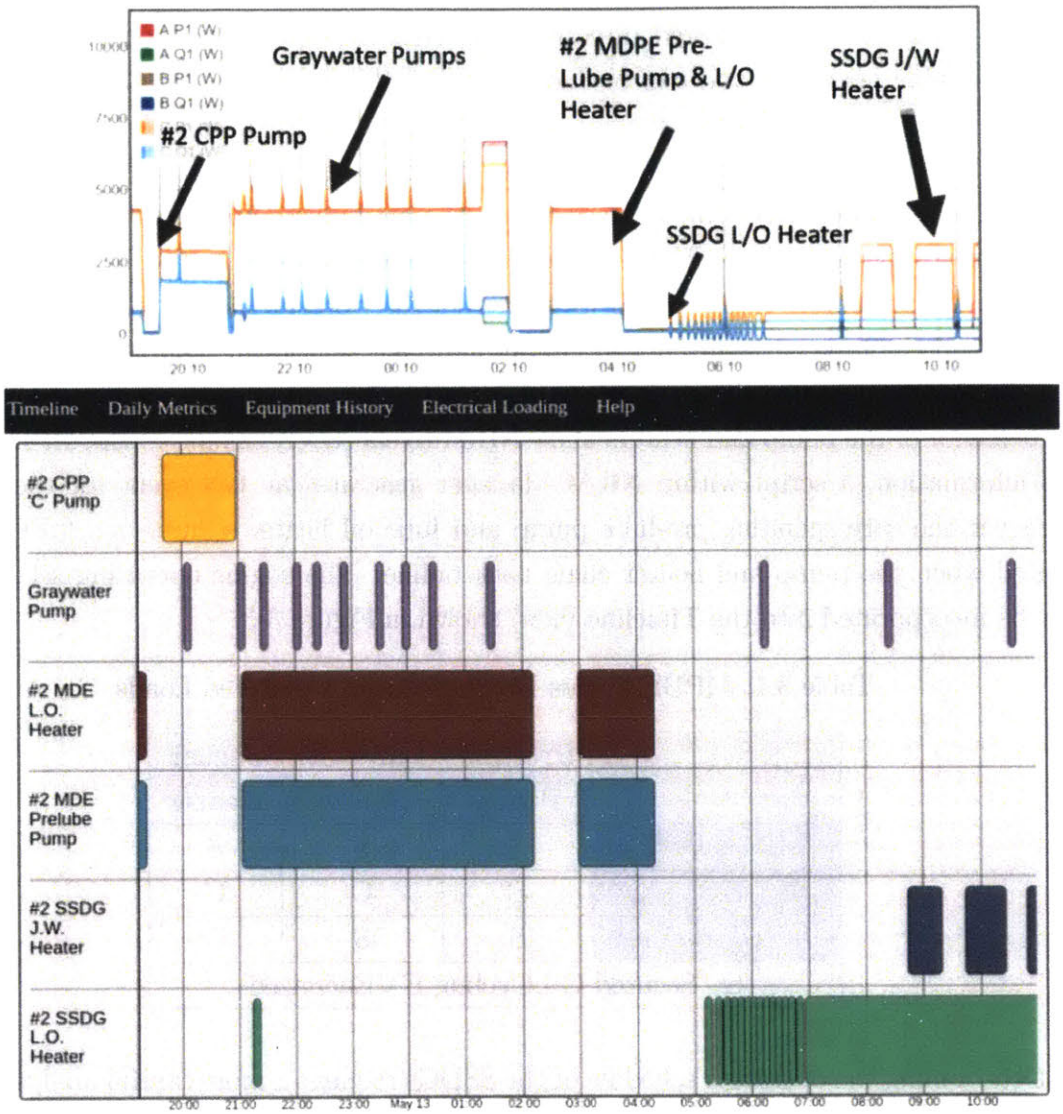

Figure 3-2: Top: Real and reactive power from 13 May 2018 on the port sub-panel on SPENCER. Bottom: NILM Dashboard Timeline view transforms power data into a schedule of equipment events that can be viewed by the crew. Colored bars represent periods where the equipment is energized.

With a brief glance at the Timeline, a watchstander or the Chief Engineer can get detailed insight into activity within the ship's engine room without the installation of large sensor network. Tracking and displaying the the load activity from the sub-panels provides the crew with insight that would have otherwise taken countless man-hours to gather and compile.

However, the behavior of select electromechanical equipment does not create a full picture of the ship's machinery. Many mission-critical systems, such as the main propulsion diesel engines (MPDEs) and ship's service diesel generators (SSDGs), do not have an electrical signature that can be monitored directly using NILM. However, the loads on the panels support engine operation and can be used to indirectly determine the status of the larger systems (see Figure 1-2).

For the MPDE, the pre-lube pump and the lube oil heater automatically secure when the engine comes on-line. When the engine secures, the pre-lube pump is activated and the lube oil heater will begin cycling. This indicates that the engine is in a standby status and can be started immediately if needed. During extended periods at the pier, the engine is fully secured and the lube oil heater and pre-lube pump are shutdown. The jacket water heater remains online to ensure the engine does not reach cold iron status. Table 3.1 shows how the behavior of the pump and heaters allow NILM to effectively "monitor" the MPDE. Using this information, a script within NILM Manager generates an ON event for the MPDEs whenever the corresponding pre-lube pump and lube oil heater secure. An OFF event is created when the pump and heater come back online. The engine operating schedule can now be incorporated into the Timeline view, shown in Figure 3-3.

Table 3.1: MPDE Status Determined by Monitored Loads

	Prelube Pump	**Lube Oil Heater**	**Jacket Heater**
MPDE Online	S	S	S
MPDE Standby	E	C	C
MPDE Secured	S	S	C
S - Secured C - Cycling E - Energized			

Determining the operating schedule of the SSDGs requires a more careful analysis of load behavior. Table 3.2 shows the SSDG status determined by the operations of the lube oil and jacket water heater. Unlike the MPDE, the heaters do not immediately energize when the engine secures, but do follow a consistent operational schedule. The period of an SSDG lube oil heater is consistently 10 minutes, with a 50 percent duty cycle. If more than 10 minutes is detected between lube oil heater operations, an ON event is created for the SSDG. When the lube oil heater comes back online, an OFF event is created and time-stamped 5 minutes into the past. The Timeline view for SSDG operations is shown in below in Figure 3-4.

Table 3.2: SSDG Status Determined by Monitored Loads

	Lube Oil Heater	Jacket Heater
SSDG Online	S	S
SSDG Standby	E	C
SSDG Secured	S	S
S - Secured C - Cycling		

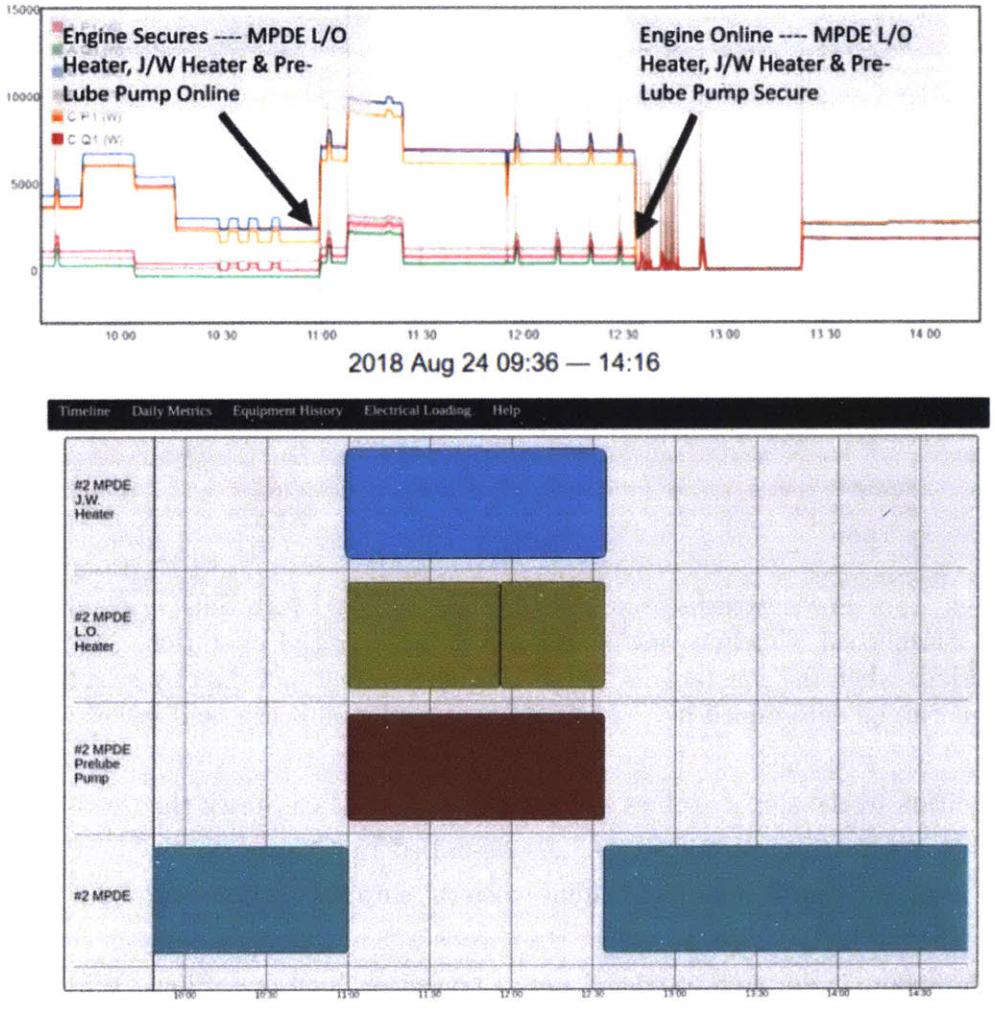

Figure 3-3: Top: Power stream from SPENCER port panel with MPDE activities tagged. Bottom: Dashboard Timeline view showing a cycle of MPDE operation. The green bar shows periods when the engine is online. The engine is not directly monitored, but its operations can be determined by the activity of the pre-lube pump, lube oil heater and jacket water heater.

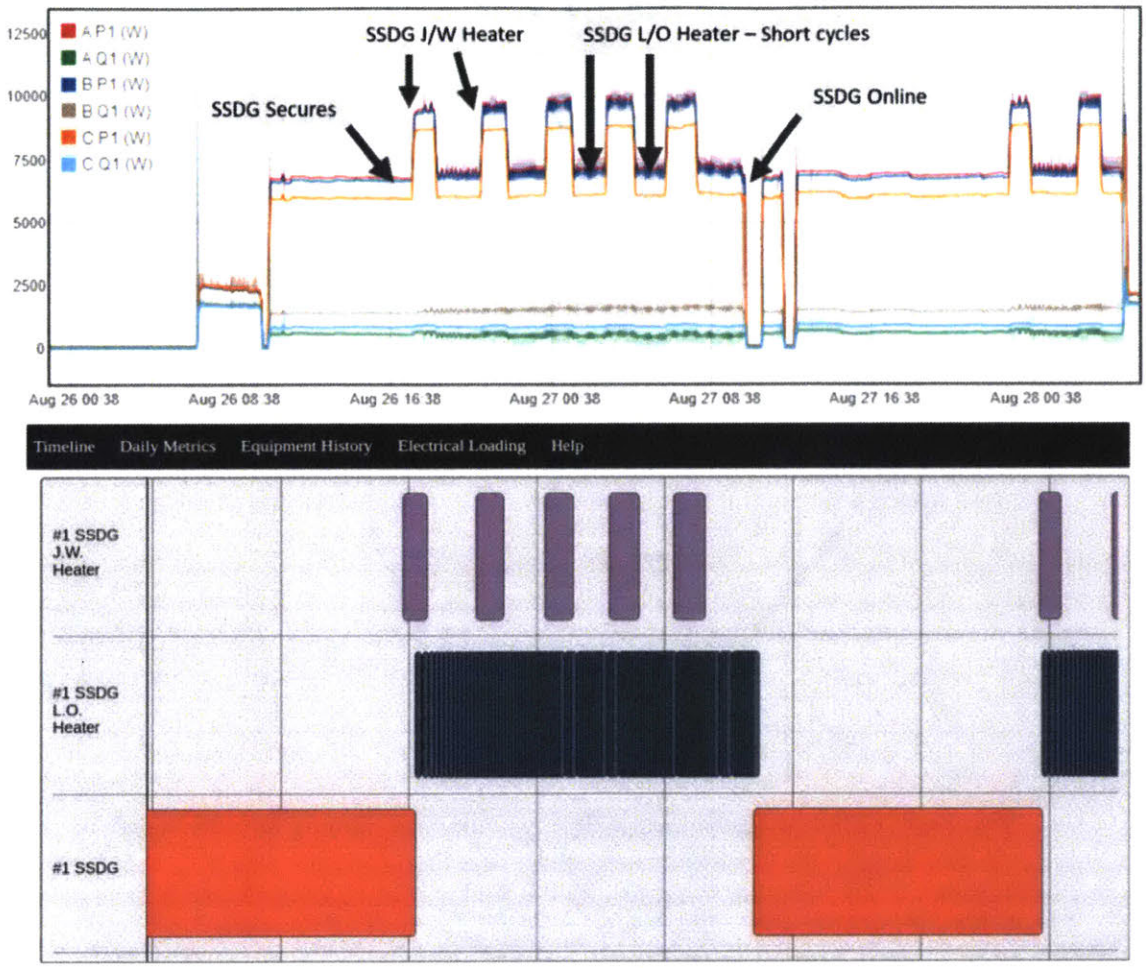

Figure 3-4: Top: Power stream from SPENCER port panel with SSDG operations tagged. Bottom: Dashboard Timeline view showing a cycle of SSDG operation. The orange bar shows periods when the generator is online. The generator is not directly monitored, but its operations can be determined by the activity of the lube oil heater and jacket water heater.

In addition to the ship's engines and generators, NILM can track the operational status of the ship itself using electric loads. As described in Chapter 2, the CPP "C" pumps supplied from the panels are brought online when the ship has set the Restricted Maneuvering Doctrine (RMD). RMD is set when the ship needs maximum redundancy in the propulsion system, for example during a restricted waters transit or flight operations. When the ship is in RMD, critical equipment in the engine room can only be secured if a class Bravo (fuel or oil) fire is imminent or with permission of the Commanding Officer. It is extremely important that engine room watchstanders are aware when the ship enters and leaves RMD. NILM can track and display RMD by monitoring the operations of the CPP pumps.

The addition of engine, generator, and RMD status allows for the creation of "Machinery

Status" interface, shown below in Figure 3-5. This view is designed to mimic the status display found in conventional machinery monitoring systems. The screen is divided into port and starboard sides showing information provided from each panel. Engines can be shown online, secured, or in standby depending on the status of their supporting equipment. Individual loads can be displayed as online, cycling, or secured. In the example below, all engines and generators are online for a restricted waters transit, and their associated support equipment is displayed as secured. Both CPP pumps are online providing additional maneuvering capability and the cutter has set RMD. The graywater pumps continuously cycle to process waste water.

The current Dashboard interface deployed on SPENCER does not support the Machinery Status view shown below. However, future iterations could adopt this screen without the additional data collection. The Dashboard is a web-based application [30, 31] that can be accessed from any computer connected on the same network as the NILM. Users on the bridge or other command and control spaces could use this information to confirm current machinery alignment in the engine room.

The remainder of this chapter describes the use of NILM as an automatic watchstander on the WMEC-270 during different operational states. Power data is used to display the Machinery Status and Timeline interfaces in order to improve situational awareness and reduce watchstander workload.

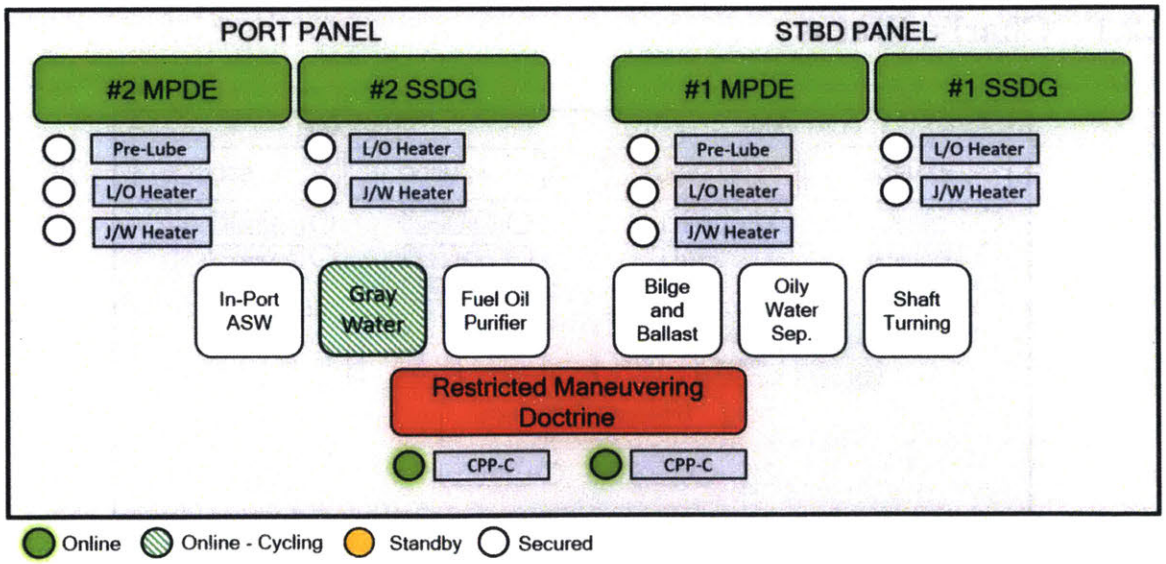

Figure 3-5: Machinery status of the WMEC-270 during a transit through restricted waters. Power data from the two monitored sub-panels can be used to track and display the operations of the ship in different operating states.

3.2 Tracking ship operations

U.S. Coast Guard policy dictates that a cutter may be in one of three operational statuses: Alpha, Bravo or Charlie [32]. During "Alpha" status, the cutter must be underway at sea. "Bravo" status indicates the cutter must be ready to enter Alpha status within a prescribed period of time. Bravo status is always accompanied with a number of hours within which the cutter is expected to enter Alpha status if the need arises. For example, a B-12 status means the cutter must be ready to get underway within 12 hours. Policy states that during Bravo status the "cutter must maintain system and equipment capabilities, be able to man watch stations and meet inport emergencies and functions to get underway within the established time period" [32]. Finally, "Charlie" status indicates that the cutter is in a maintenance period or otherwise not available for employment.

Machinery will have different behavior depending on the operational status of the cutter. Systems are progressively brought from secured status into standby and eventually online as the cutter transitions from Charlie to Alpha. Furthermore, during Alpha status different equipment may be online to support special maneuvering (i.e. RMD) or other operations. Using the Timeline view on the NILM Dashboard and the information from the Machinery Status screen proposed above, cutter crews can confirm that the behavior within the engine room matches the operational state of the ship.

3.2.1 Charlie Status

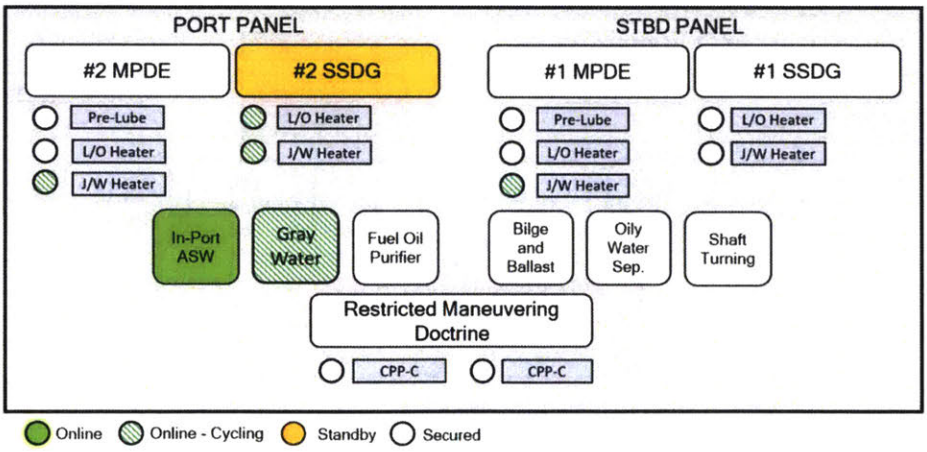

Figure 3-6: Machinery status created for a Charlie period using NILM data from two sub-panels. Power is provided from shore and one SSDG is in standby.

The cutter is in "Charlie" status when moored at the pier for extended periods of time, normally in home port. During these periods, NILM data should show a quiet engine room. The main diesel propulsion engines should be secured and kept warm using only the jacket water heaters. The cutter is receiving power from a shore tie connection, but a minimum of one SSDG must be kept in standby status in the event of a power interruption. The standby generator's jacket water and lube oil heater should be cycling. Cooling water to the ship is provided by the inport ASW pump and the gray water system remains active for the duty crew. Normal inport Charlie operations are shown in Figure 3-6 and Figure 3-7.

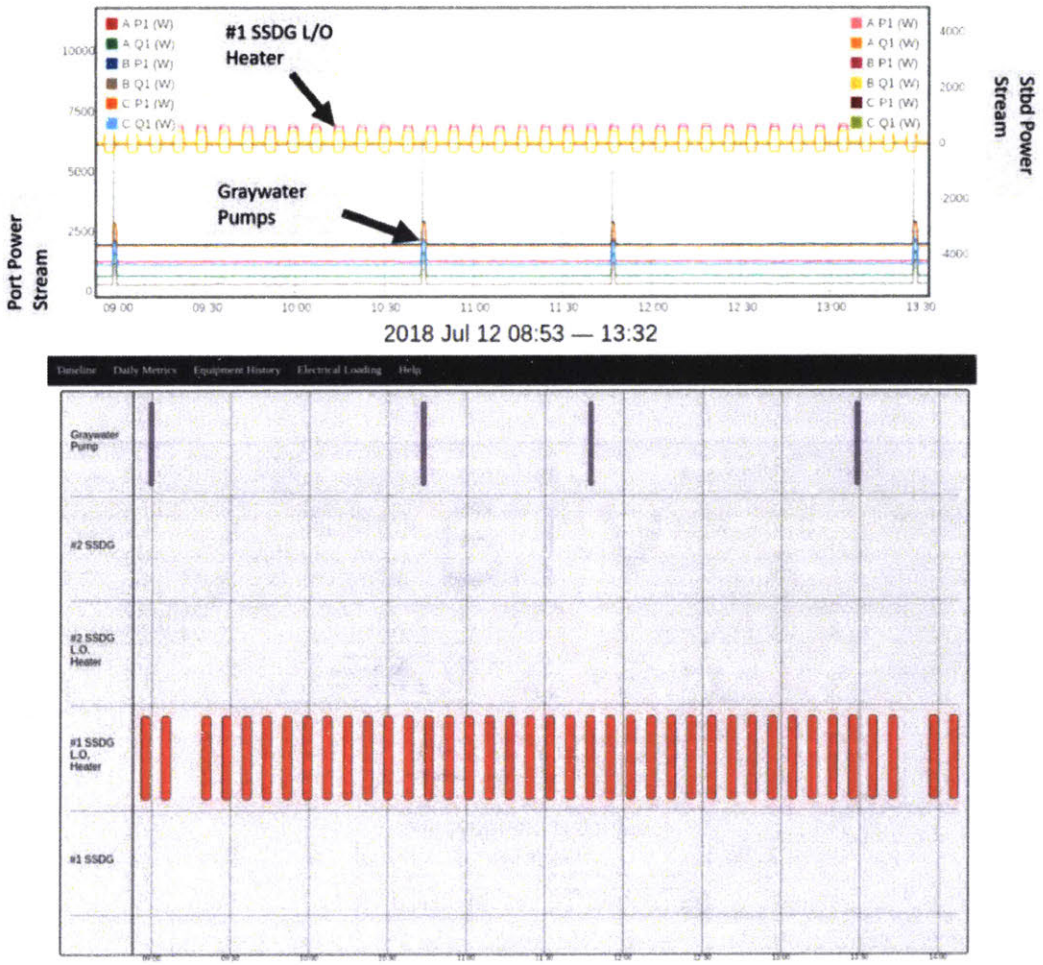

Figure 3-7: Top: Port panel power stream (lower) and Stbd panel power stream (upper). Bottom: NILM Dashboard Timeline from a quiet day in home port Charlie status.

Non-intrusive load monitoring can alert inport Charlie watchstanders to improper operations. It is essential that at least one SSDG be in stand-by status at all times to prevent a "darkened ship" state in the event of a shore-side power failure. The oily water separator

and CPP pumps should rarely if ever be online while in Charlie status, unless maintenance is being conducted on those systems.

3.2.2 Bravo - 2 Hour Standby

The operations of engine room equipment drastically change as the cutter transitions to a Bravo status from inport Charlie. Figure 3-9 shows the Machinery alignment for Bravo-2 status. The pre-lube pumps and lube oil heaters are brought online prior to MPDE start and both SSDG are brought to standby status. The inport ASW pump is secured as salt water cooling is shifted to the at-sea configuration. The CPP "C" pumps may be brought online to test the propulsors before the cutter leaves the pier. The Machinery Status interface below in Figure 3-8 shows both MPDEs and both SSDGs have shifted to standby status.

In a Bravo-2 state, watchstanders can check the NILM Machinery Status or Timeline views to ensure that MPDE and SSDG pumps and heaters are functioning and that the cutter is ready to get underway. As described in a case study the next Chapter, automatic controllers can display that engine heaters are online even when a failure has occurred, impacting engine readiness. Non-intrusive load monitoring can detect these faults and bring them to the attention of the ship's crew.

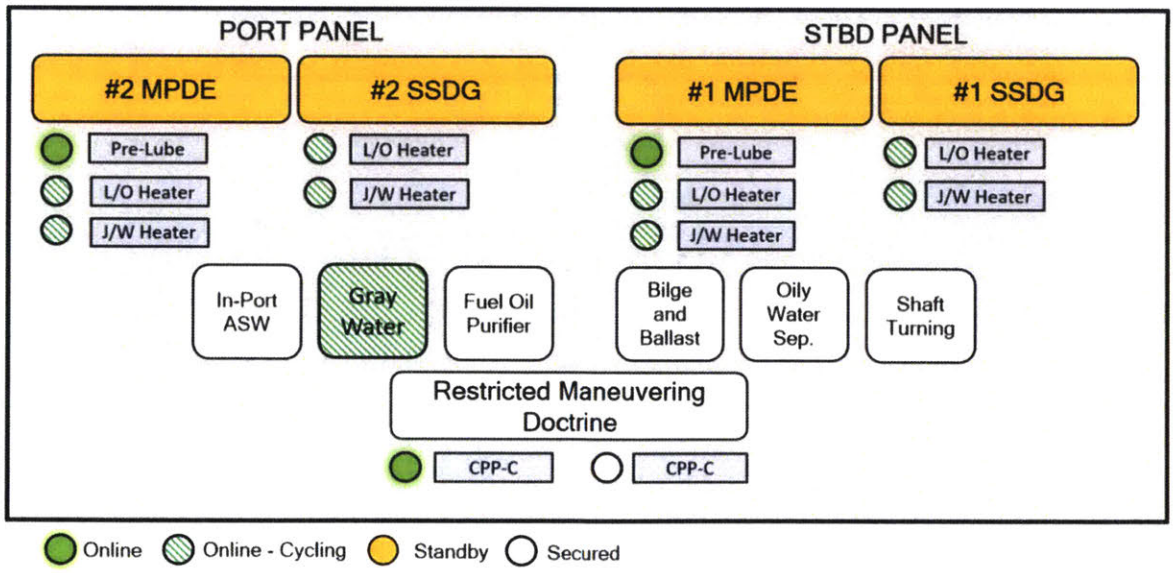

Figure 3-8: Machinery status created for a Bravo-2 period using NILM data from two sub-panels. All engines and generators are in standby status in preparation to get underway.

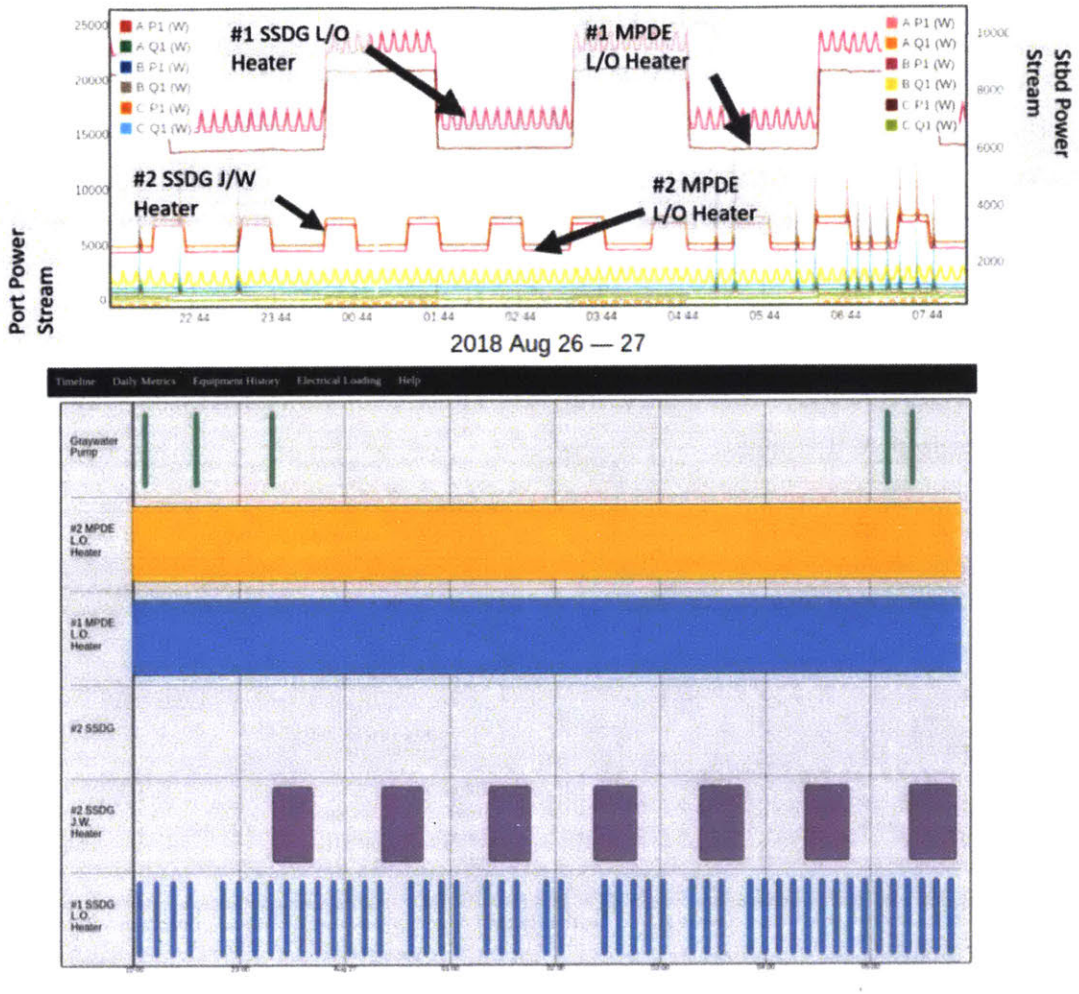

Figure 3-9: Top: Port panel power stream (lower) and Stbd panel power stream (upper). Bottom: NILM Dashboard Timeline from a period of Bravo-2 status. Engine heaters are active for both MPDE and both SSDG in preparation to get underway.

3.2.3 Alpha Status - Underway RMD

When the cutter is in Alpha status, the machinery can be aligned in a variety of different configurations depending on the ship's mission. Immediately upon exiting Bravo status and getting underway, the cutter will be in an Underway-RMD status due to the close proximity of other vessels and navigational hazards. All engines and the generators must be online and their respective heaters should be secured. Both CPP "C" pumps should be operating, providing the redundancy required during RMD. Auxiliary systems such as the Fuel Oil Purifier (FOP) are normally secured while the ship is in restricted waters. During underway RMD, any deviation from the Machinery Status or Timeline views shown below may indicate improper operation of the plant.

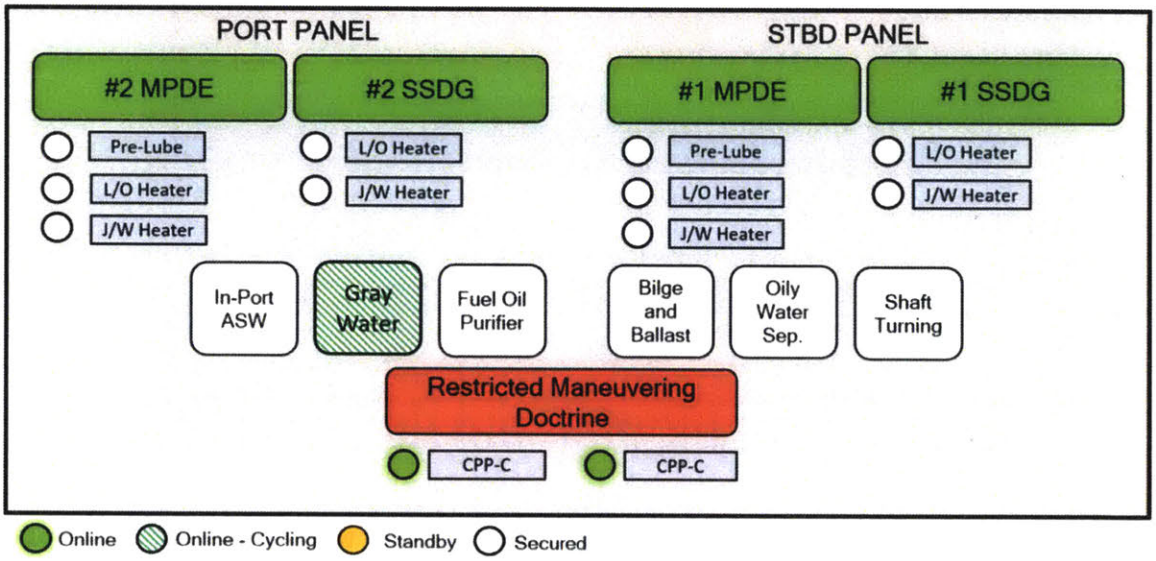

Figure 3-10: NILM Dashboard Timeline of equipment operations as the cutter transitions from Charlie status to underway RMD status. All engines and CPP pumps are brought online.

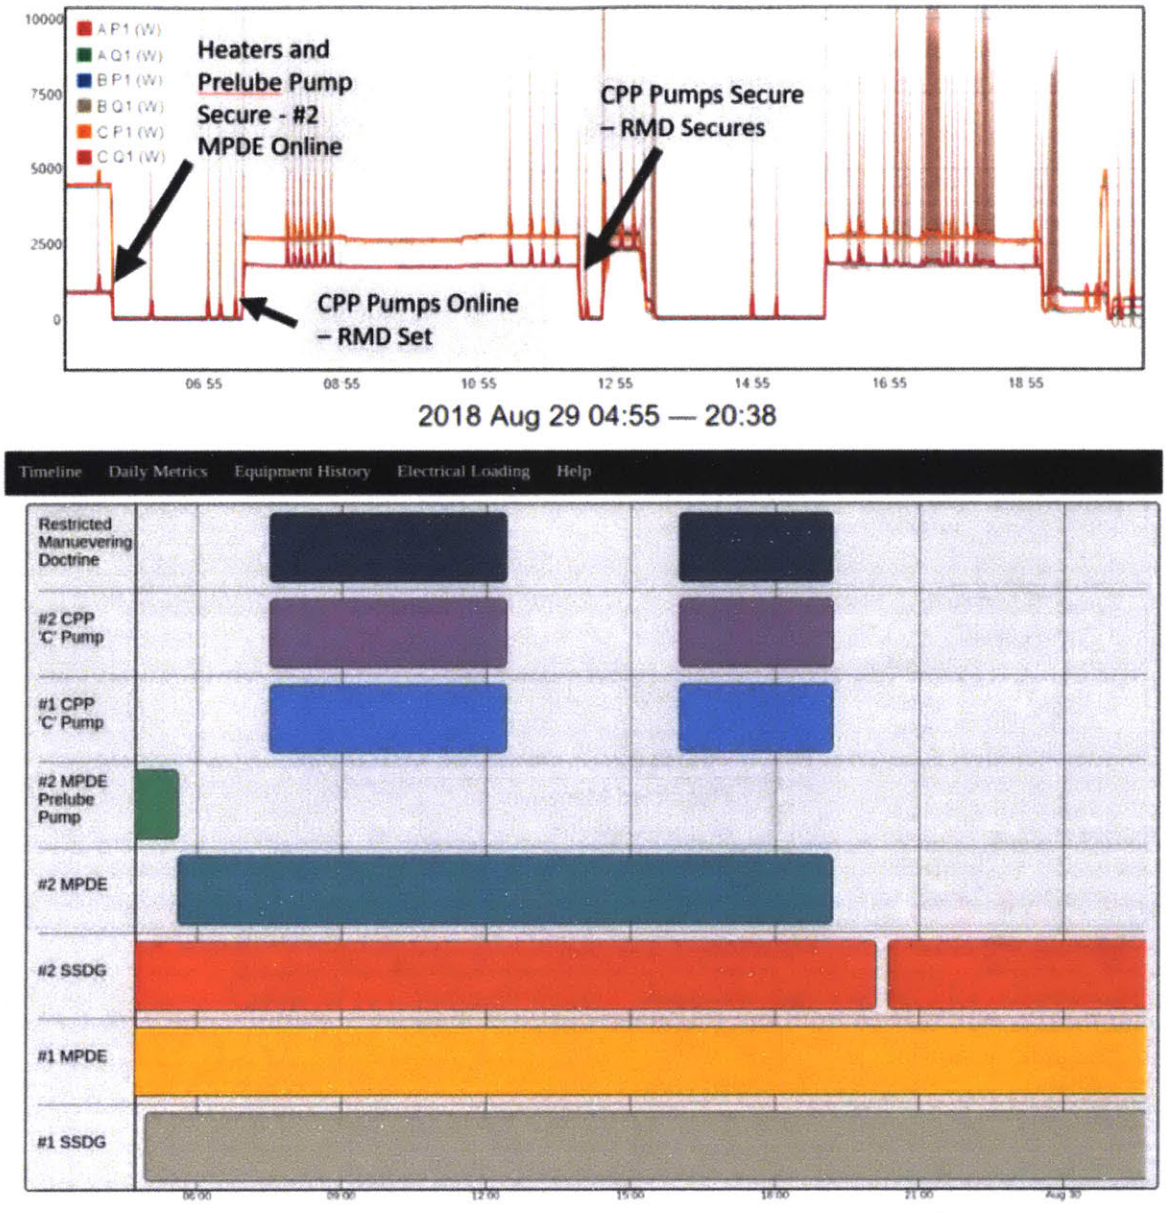

Figure 3-11: Top: Port panel power stream as the cutter gets underway from the pier. Bottom: NILM Dashboard Timeline as the cutter gets underway. All engines are shown online and the CPP "C" pumps are energized for increased maneuvering capability.

3.2.4 Alpha Status - Normal Transit

Once the cutter leaves restricted waters and enters the open ocean, a more standard Alpha alignment in the engine room can be observed by NILM. The CPP "C" pumps are secured and cutter stands down from RMD. One engine and one generator are secured and placed in standby status to save fuel, indicated by the associated heaters and pumps cycling. The fuel oil purifier and oily water separator may be online to move fuel and waste to appropriate tanks. Non-intrusive load monitoring data during an at sea transit are shown in Figure 3-12 and Figure 3-13.

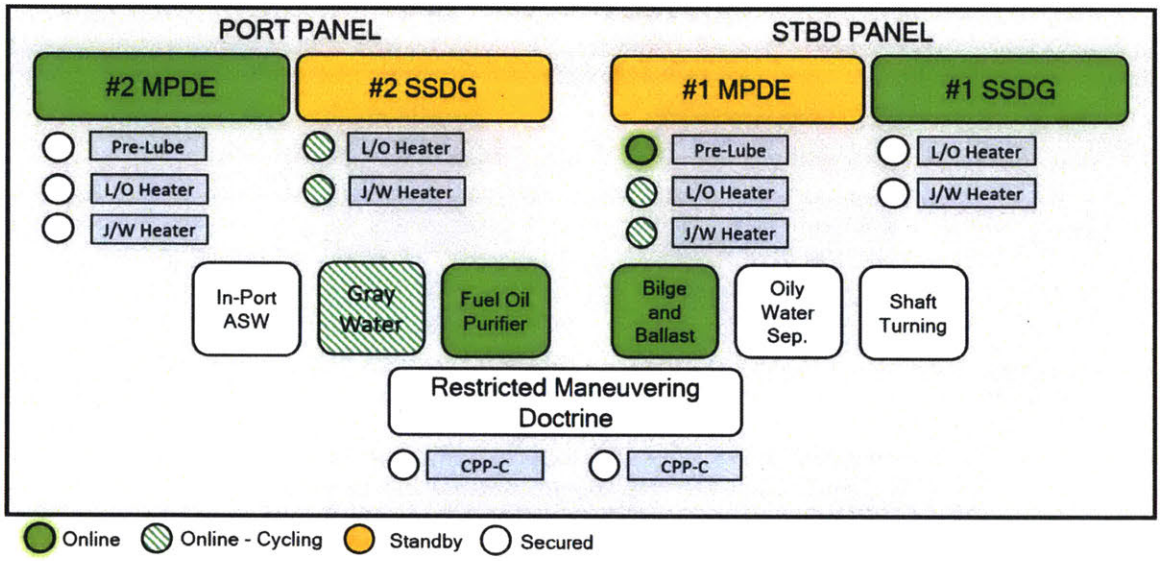

Figure 3-12: NILM data from the subpanels captures machinery status as the crew transits in the open ocean.

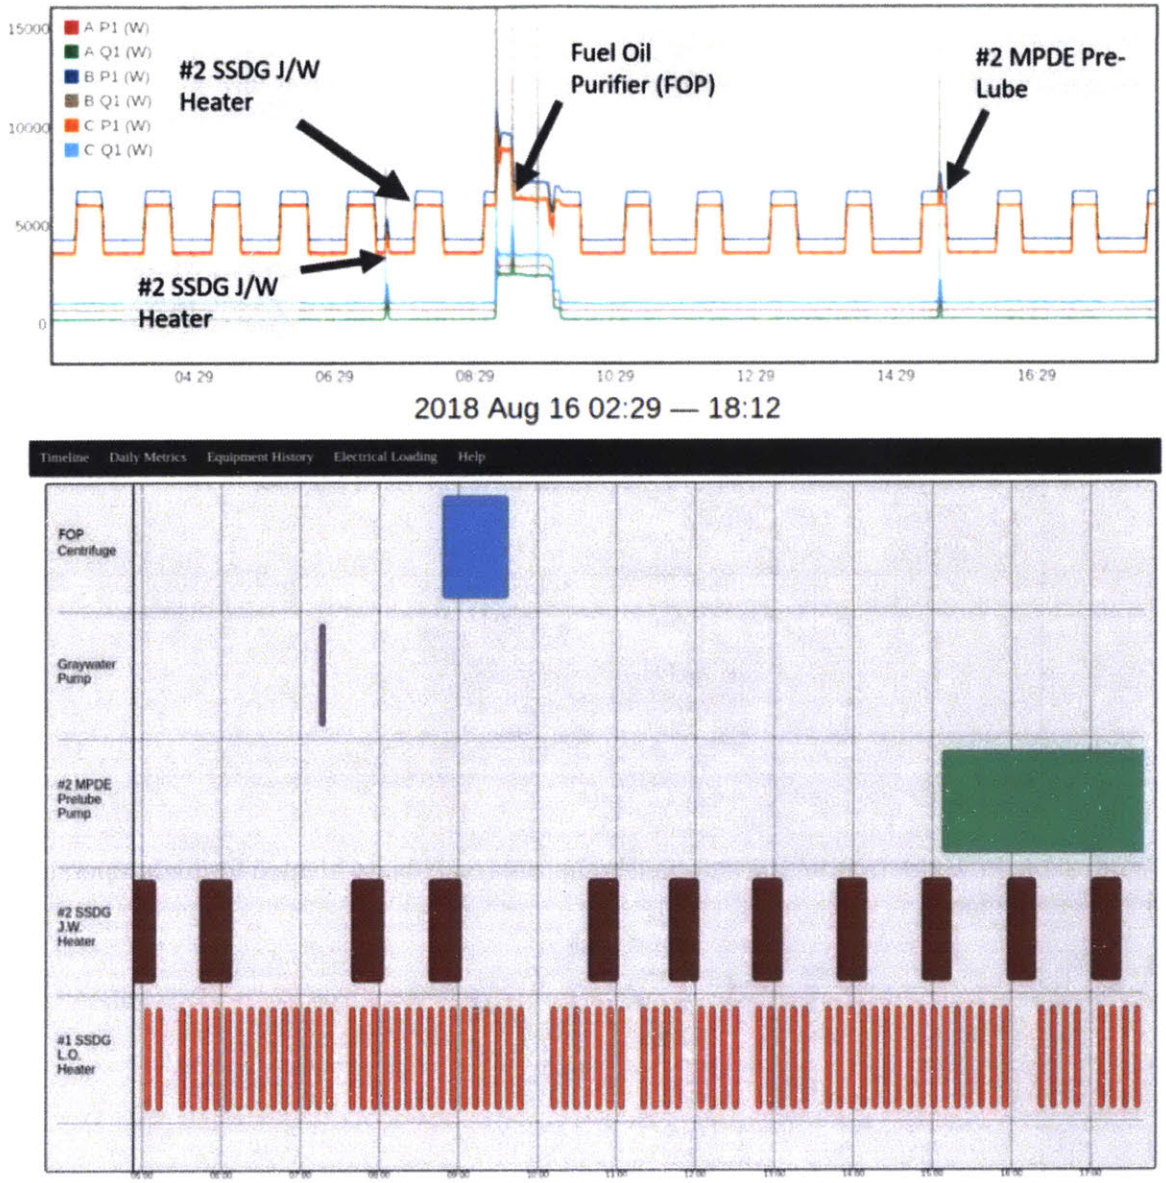

Figure 3-13: Top: Port panel power stream during an underway transit. Bottom: Dashboard Timeline view during Alpha status. The Fuel Oil Purifier is activated briefly to refill the fuel oil service tanks and one SSDG is in standby status to manage generator loading.

3.2.5 Alpha - Underway Operations

During a busy day of underway operations, the crew will energize and secure engines throughout the day. The cutter may also set RMD to conduct law enforcement or flight operations, indicated to NILM by the CPP pumps energizing and securing. The cutter crew may opt to take on ballast water via the Bilge and Emergency Ballast (BEB) pump to improve ship stability in the open ocean. Typical underway operations captured by NILM are shown in Figure 3-14 and Figure 3-15.

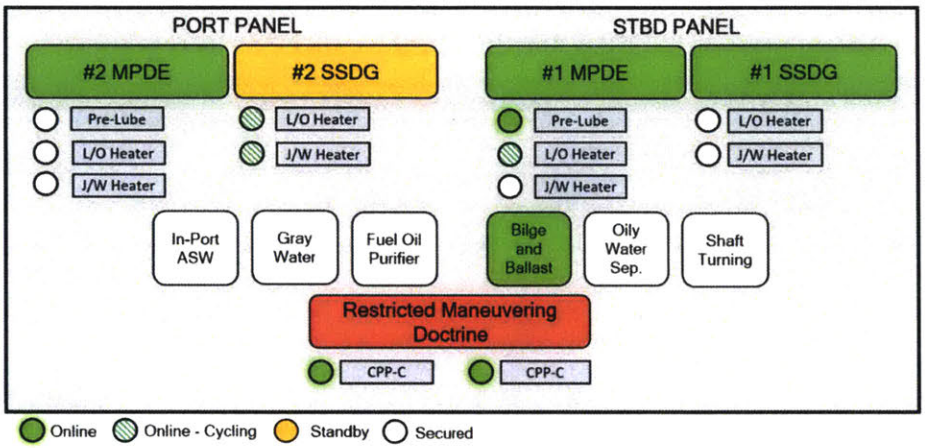

Figure 3-14: Power data from the sub-panels captures machinery aligned for flight operations or law enforcement.

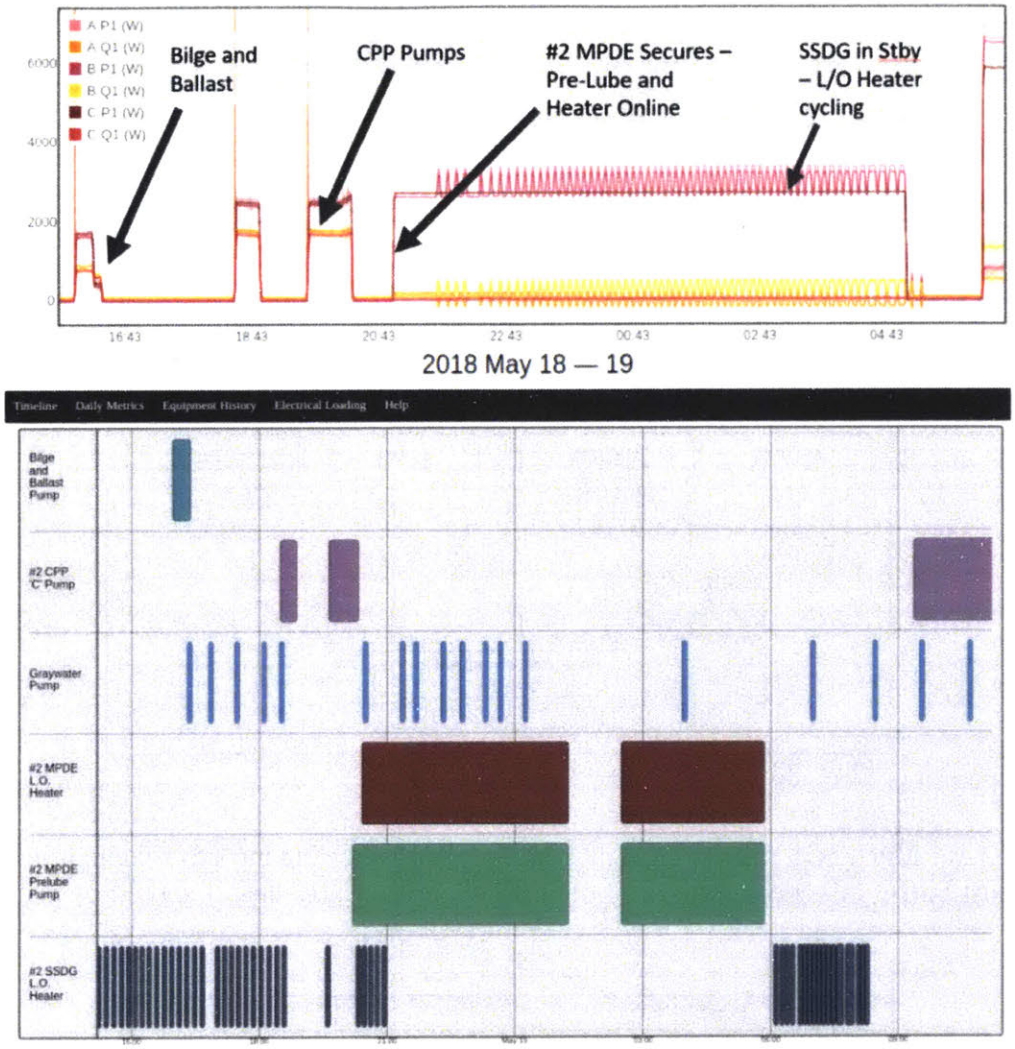

Figure 3-15: Top: Port panel power stream during underway operations. Bottom: Dashboard Timeline view during operations at sea. CPP "C" pumps are actuated while the cutter maneuvers conducting the mission. Engines and generators are brought online and secured as needed.

3.2.6 Incorrect Machinery States

Non-intrusive load monitoring can quickly draw a watchstander's attention to improper machinery configurations within the engine room. In Figure 3-16 below, the cutter is in a Charlie status, with the starboard SSDG in standby. However, the starboard SSDG L/O heater is not cycling, meaning that the engine would not be ready for operation in the event of an emergency. Additionally, the port MDPE jacket water heater is not operating, leaving the engine in cold iron status. This may be due to maintenance, but could also be due to a

fault condition. A watchstander could with access to NILM information could investigate and troubleshoot these issues before the cutter needs to shift to a more active status. Power data can serve a similar function while underway. In Figure 3-17, the cutter is in Alpha status, with the port MPDE online and the starboard MPDE in standby. However, the starboard MPDE jacket water heater is not running, an indication of a possible fault. Additionally, the port CPP "C" pump is operating but the starboard CPP "C" pump is offline, indicating a potential loss of pressure in the port CPP system. The examples below clearly demonstrate how NILM can improve watchstander situational awareness in any cutter status.

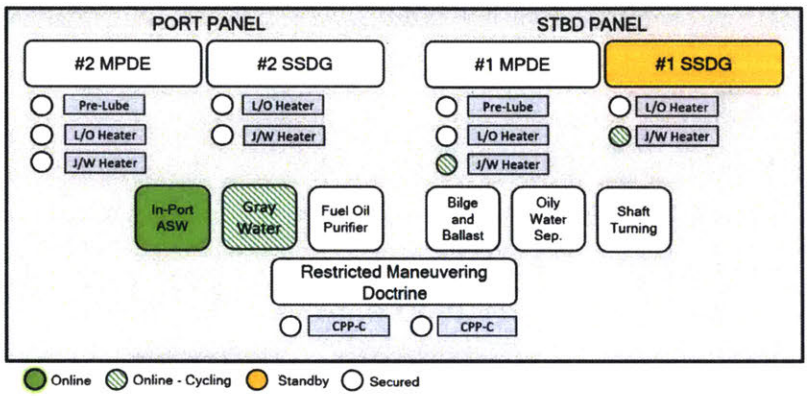

Figure 3-16: Machinery status during a Charlie period showing incorrect configuration. The starboard SSDG lube oil heater and port MPDE jacket water heater should be cycling.

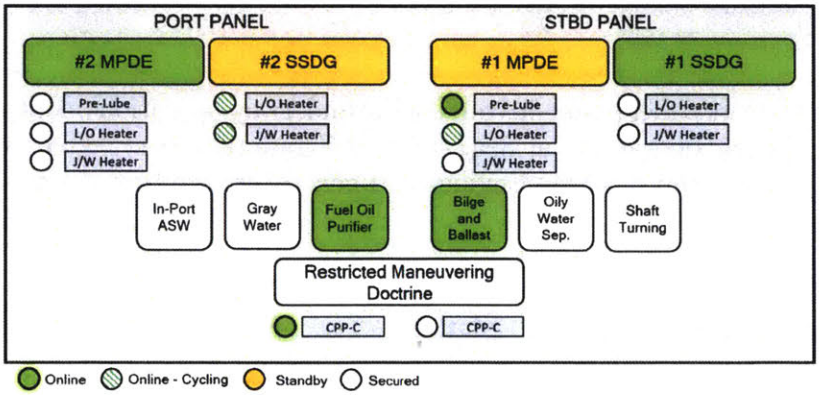

Figure 3-17: Machinery status during an Alpha period showing incorrect configuration. A single CPP pump is online, indicating a non-standard propulsion configuration.

Chapter 4

Fault Detection Case Studies

Automation and closed-loop feedback control can spell disaster for efforts to perform condition based maintenance and energy economization. As the performance of a critical component in a system degrades, an automatic controller will compensate to maintain commanded output levels. By design, feedback control works to mask the effect of "soft faults," variations in internal component performance in a larger system that do not fully stop the system. The system will continue to operate and provide apparently acceptable service while an automatic controller forces more energy consumption and induced wear. Soft faults may elude even careful watchstanders. These situations can persist for expensively long periods before a "hard fault" finally stops the system, alerting operators but also creating a casualty that may cripple mission readiness.

Systems for heating, ventilation, and air conditioning (HVAC) provide telling examples. References [33] and [34] describe surveys of faults in commercial HVAC and roof-top package units for environmental control. Reference [33] reviews the aggregate results of four field surveys that examined 503 roof-top HVAC units in 5 states at 181 different commercial building sites. Approximately half of the surveyed units were hobbled with soft faults that did not prevent system operation. For example, reference [33] reports that 46% of the units had improper refrigerant charge, 64% of the units had problems with their economizers, and 42% of the units had blockages or other interference with airflow. Repairs to these problems brought significant energy savings, e.g., 14 to 40 percent in the case of economizer repairs.

Reference [34] notes the difficulty of finding soft faults in comparison to hard faults, while highlighting the significance of soft faults in affecting energy consumption and induced wear. Enhanced vigilance and testing can find soft faults, but at a price [34]. For mission critical systems, enhanced inspection leads to increased monitoring effort and an expanded burden for operators [35]. Enhanced sensing and inspection also potentially leads to a greater reliance on the accuracy and function of a larger number of sensors, creating new

fault points in the system. Sensors can record a huge collection of process information, including speed, temperature, pressure, operating time, light levels, sound levels, viscosity, gas content, and so on. Poorly applied networked sensing solutions simply add to the burden of maintaining a system by making a raft of raw data available to a busy watchstander. Efforts to apply machine learning or automatic analysis of this data, another contemporary trend, are often unsatisfying and difficult to trust when used to search for correlations in data without consideration of the underlying physics [35].

Feedback control cannot hide the changes in electrical energy consumed by compromised system components. NILM provides a flexible, low-cost sensing platform that can detect soft faults through careful observations of electric power data. Furthermore, rather than adding another sensor and set of data for operators, NILM has the capability to analyze data in real-time and provide actionable recommendations through the NILM Dashboard (Appendix B) platform. This chapter presents a series of soft-faults detected by NILM on the WMEC-270 and discusses the effectiveness of the NILM Dashboard in detecting and communicating these faults to watchstanders.

4.1 Diesel Engine Jacket Water Heater Failure

Described earlier in Chapter 2, the WMEC-270 MPDE jacket water heater consists of two 4.5kW heating elements on either side of the engine block, shown below in Figure 4-1. The heaters appear to a NILM as a single electrical load, 9kW total and 3kW per phase. The controller configuration causes the MPDE jacket water heater and lube oil heater to operate together while the cutter is underway. Figure 4-2 shows the electrical transients for a healthy MPDE jacket water heater operating alone and with the MPDE lube oil heater.

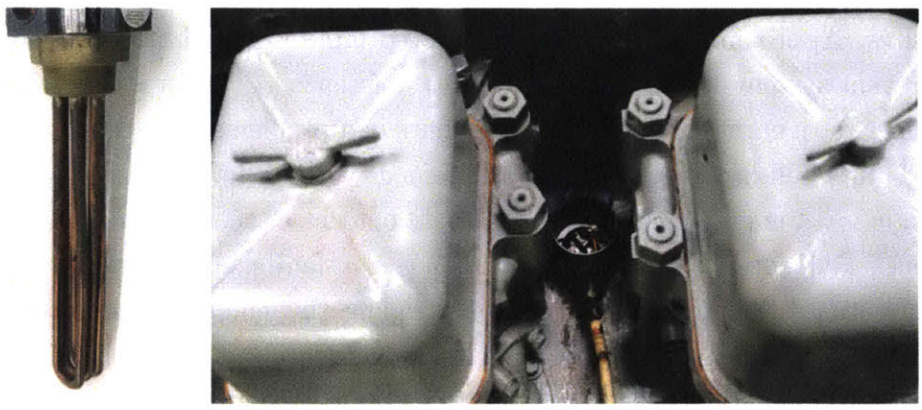

Figure 4-1: Left: New MPDE jacket water heater. Right: Jacket water heater mounts to the engine block between the cylinders.

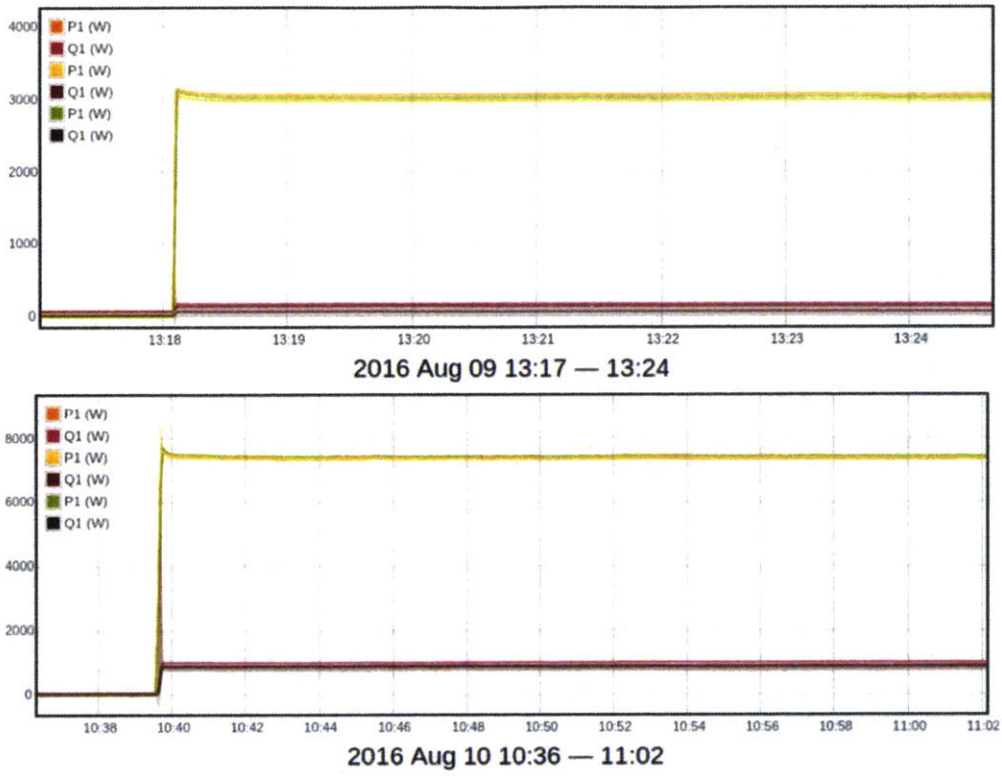

Figure 4-2: Top: ON transient for a healthy jacket water heater from SPENCER starboard MPDE, showing 3kW per phase. Right: When the ship is in Alpha status, the jacket water heater and lube oil heater operate together as a 7kW resistive load.

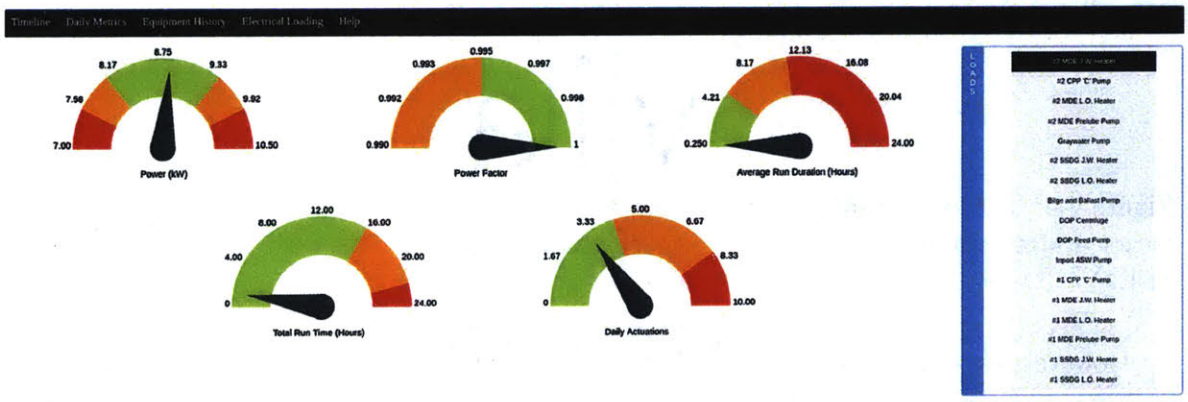

Figure 4-3: NILM Dashboard showing a healthy MDPE jacket water heater condition. All gauges are within the green regions.

On the NILM Dashboard, healthy heater operation is captured on the Metrics gauges, shown in Figure 4-3. 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.

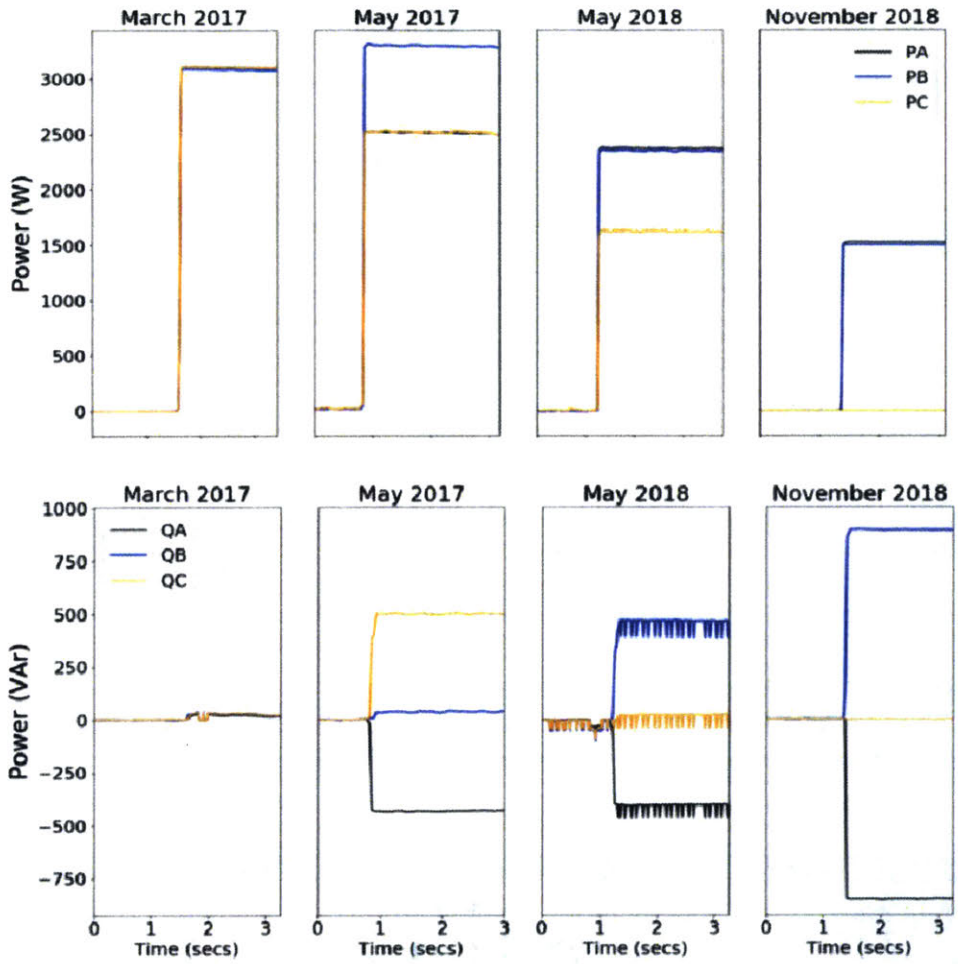

Figure 4-4: Turn-ON transients in real power (top) and reactive power (bottom) for MPDE jacket water heaters, observed over more than a year on the port engine on USCGC SPENCER.

However, as seen in Figure 4-4, 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 from 3000W to 1500W per phase. During the same period, the heater, normally a purely resistive load, showed a varying reactive component 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 4-4 must be properly identified and tracked. The artificial neural network (ANN) approach used to identify loads on the SPENCER and ESCANABA (described in Appendix B), relies on supervised learning where load transients are hand-labeled in order to train the identification algorithm [36]. 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. This will reduce the fault detection capabilities of the NILM Dashboard.

Despite the challenges presented by the changing load signature, the NILM Dashboard was still able to effectively capture this case with no adjustment to the identification algorithm. Figure 4-5 shows the gauge tracking real power moving from the green region into the yellow at the first sign of abnormality (May 2017 in Figure 4-4). This indicator should prompt a watchstander to check to inspect the heaters. Once the heater reaches the state shown in May 2018 in Figure 4-4, the real power draw has decreased enough that the gauge enters the red region, indicating that a failure has likely occurred. The Dashboard is effective in diagnosing this fault for two reasons; first, the combined actuation of the jacket water heater and lube oil heater reduced the effect of the diminished power demand on the classification algorithm. Secondly, there were no other loads on the panels that closely matched the features of the degraded MPDE heater transients.

The nature of this fault can be applied to improve future fault detection. The ANN neural net can be trained to identify the transients in Figure 4-4 as jacket water 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 the importance of having a parameter that can detect the lack of operation as a fault. In the event a transient is no longer recognizable, the "Runs Per Day" metric should capture the fault detection by dropping to zero.

Another strategy to improve fault detection would be to utilize the interdependence of ship loads. Equipment on the ship often functions as a finite state machine (FSM) [15], where one load should not be detected without the presence of another. In this case, if a lube oil heater activation is repeatedly detected without a jacket water heater activation, this should be displayed as a fault condition. The interdependence of loads can help ensure load identification algorithms do not miss fault conditions where the shape electrical signature is affected.

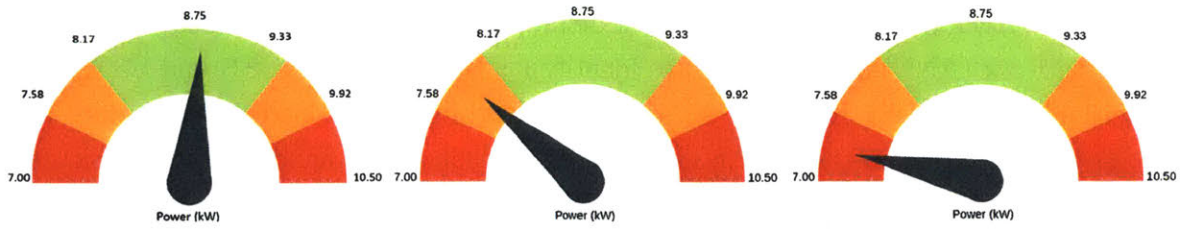

Figure 4-5: The Real Power gauge from the NILM Dashboard captures the degradation of an MDPE jacket water heater. Left: Healthy state. Center: Yellow region indicates possible fault condition. This corresponds to the May 2017 heater transient shown above. Right: A fault is present, with multiple phases now offline due to corrosion. This corresponds to the May 2018 electrical transient shown above.

The electrical signature degradation in Figure 4-2 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 4-6 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 4-7). 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 in the heater shown in Figure 4-7, 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. 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 addition of the NILM Dashboard into the engine room allows a watchstander to recognize subtle changes without altering the watch routine.

The root cause of the corrosion is unknown at this time, but it is clear that this is a systemic issue. Of the 8 heaters in place on ESCANABA and SPENCER, 7 were observed to have suffered complete failure or were operating with reduced heating capability. A memorandum regarding this failure was sent to US Coast Guard MECPL (Medium Endurance Cutter Product Line) for increased visibility.

Figure 4-6: MPDE Jacket Water heaters removed from engines after electrical degradation was observed. Top: SPENCER port inboard heater, removed December 2018. Bottom: ESCANABA port Inboard heater, removed April 2019.

Figure 4-7: Left: Bare wiring found in SPENCER port inboard jacket water heater. Right: Holes (circled) in the heating element that allowed water to enter electrical enclosure.

4.2 Graywater System Fault

The graywater system on the WMEC-270 continuously collects and processes waste water from showers and sinks. Two pumps, monitored by NILM, alternate runs to empty the collection tank either over the side or to a large holding tank, depending on the ship's position. The pumps are controlled by high and low level sensors in the tank, and should run for approximately 60 seconds, 1-2 times per hour (see Chapter 2 for additional detail). Figure 4-8 shows a healthy electrical signature and the typical operational behavior for the graywater pumps, and Figure 4-9 shows the Dashboard Metrics for a normal day of graywater operations.

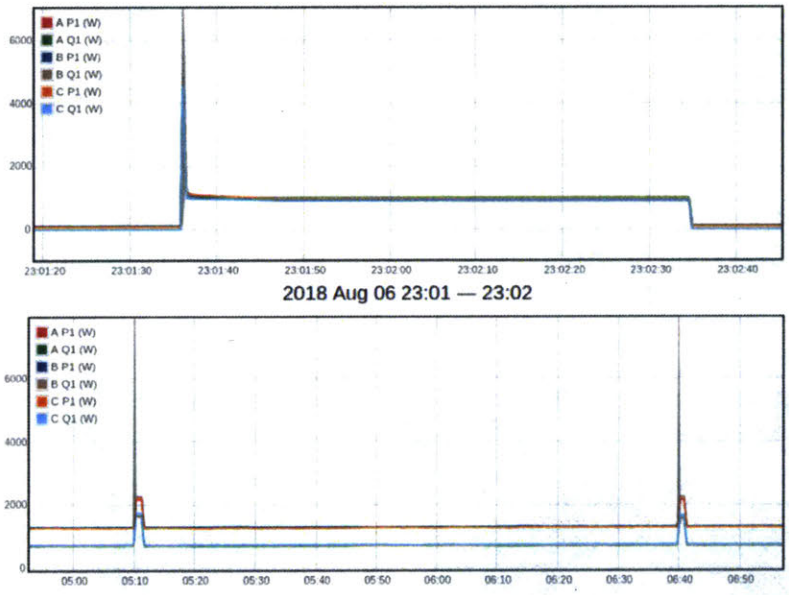

Figure 4-8: Top: Normal graywater pump electrical signature. Pump runs for approximately 60 seconds. Bottom: Normal graywater pump behavior. Pumps run on average every 1.5 hours.

During one underway cruise, the high level sensor in the graywater system became fouled, giving an electrical signal of "high level" in the tank at all times. This caused the graywater pump to run in shortened cycles, actuating for 10 seconds, 1-2 times per minute. The electrical signature and pump behavior for this fault condition is shown in Figure 4-10. A comparison between the normal condition in Figure 4-8 and the fault condition in Figure 4-10 shows that although a fault is present, the electrical signature of pump remains the same. This demonstrates an important distinction between the corrosion of the heaters described in the last section and the sensor failure shown here; for a sensor fault, the suspect equipment

remains easily identifiable.

When considering modeling failures in pumps versus heaters, it important to note that material wear in a pump does not normally manifest as the large changes in real or reactive power demand. Wear will likely characterized as gradual increases in run duration or energy consumption [13, 14]. There is likely no need to train the ANN to find a decaying pump transient, unless a specific failure mode has been identified, as seen in the next section. Gradual changes in electrical signature or sensor failures that alter equipment behavior can be detected without any modification to the load identification algorithm.

Figure 4-9: NILM Dashboard Metrics for a day of normal graywater pump operations.

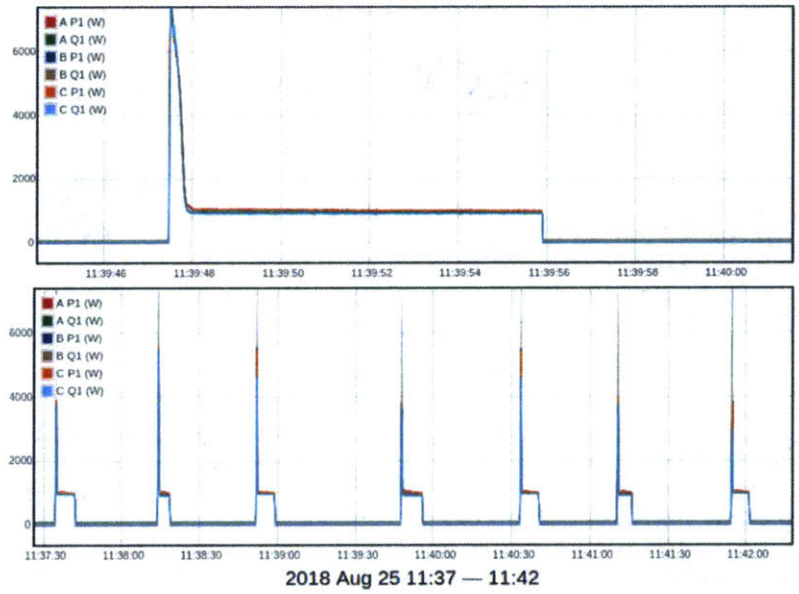

Figure 4-10: Top: Graywater pump electrical signature during sensor failure. Steady-state real and reactive power demand are identical to healthy operations. Bottom: Graywater pump behavior during sensor failure. Pump runs approximately once every minute.

Knowing that load identification will be likely be successful in a fault condition, we can theorize about the effect various fault conditions will have on pump behavior. Broken tank level sensors will change the run duration of the pump and the time between pump runs. If sinks or showers on board the ships are left running, the time between graywater pump runs will decrease. Mechanical wear on the pumps may cause them to take longer to drain the tank. Figure 4-11 depicts how various types of faults can effect the operational behavior of the graywater system. Examination of disaggregated graywater pump run data from the data revealed the plot shown at right in Figure 4-11, a clear indicator of a high level sensor failure.

Sensor failures are immediately captured on all interfaces of the NILM Dashboard, shown in Figure 4-12. On the day that the level sensor failed, the gauges for total daily run time and daily actuations move into the red region. The historic view shows a clear trend of increasing daily actuations after a steady baseline of around 25 runs per day. Furthermore, a quick glance at the Timeline for the graywater pump displays betrays an abnormal state.

Sensor failures of this nature are common and very difficult for crewmembers to detect. The system continued to process gray water as designed, but at the cost of greatly increased pump wear and energy consumption during the time of the fault condition [14]. Non-intrusive load monitoring can detect such sensing failures nearly immediately and allow crew members to address them before equipment failure occurs.

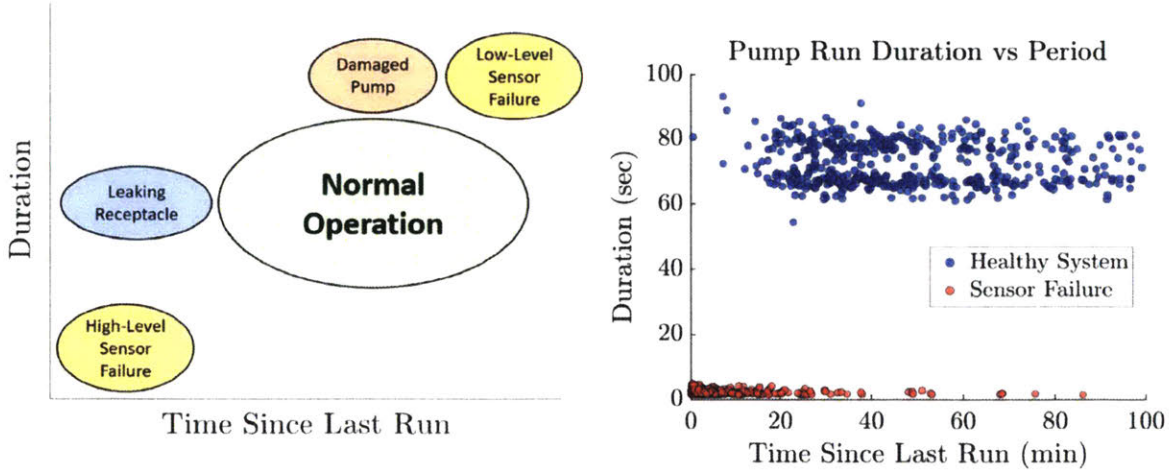

Figure 4-11: Left: Theoretical regions of normal operation and operation under various fault conditions for the gray water system, plotted on a graph of pump run duration versus time since last pump run. Right: Scatter plot of pump run data from USCGC SPENCER before and after a high-level sensor fault in the gray water system [1].

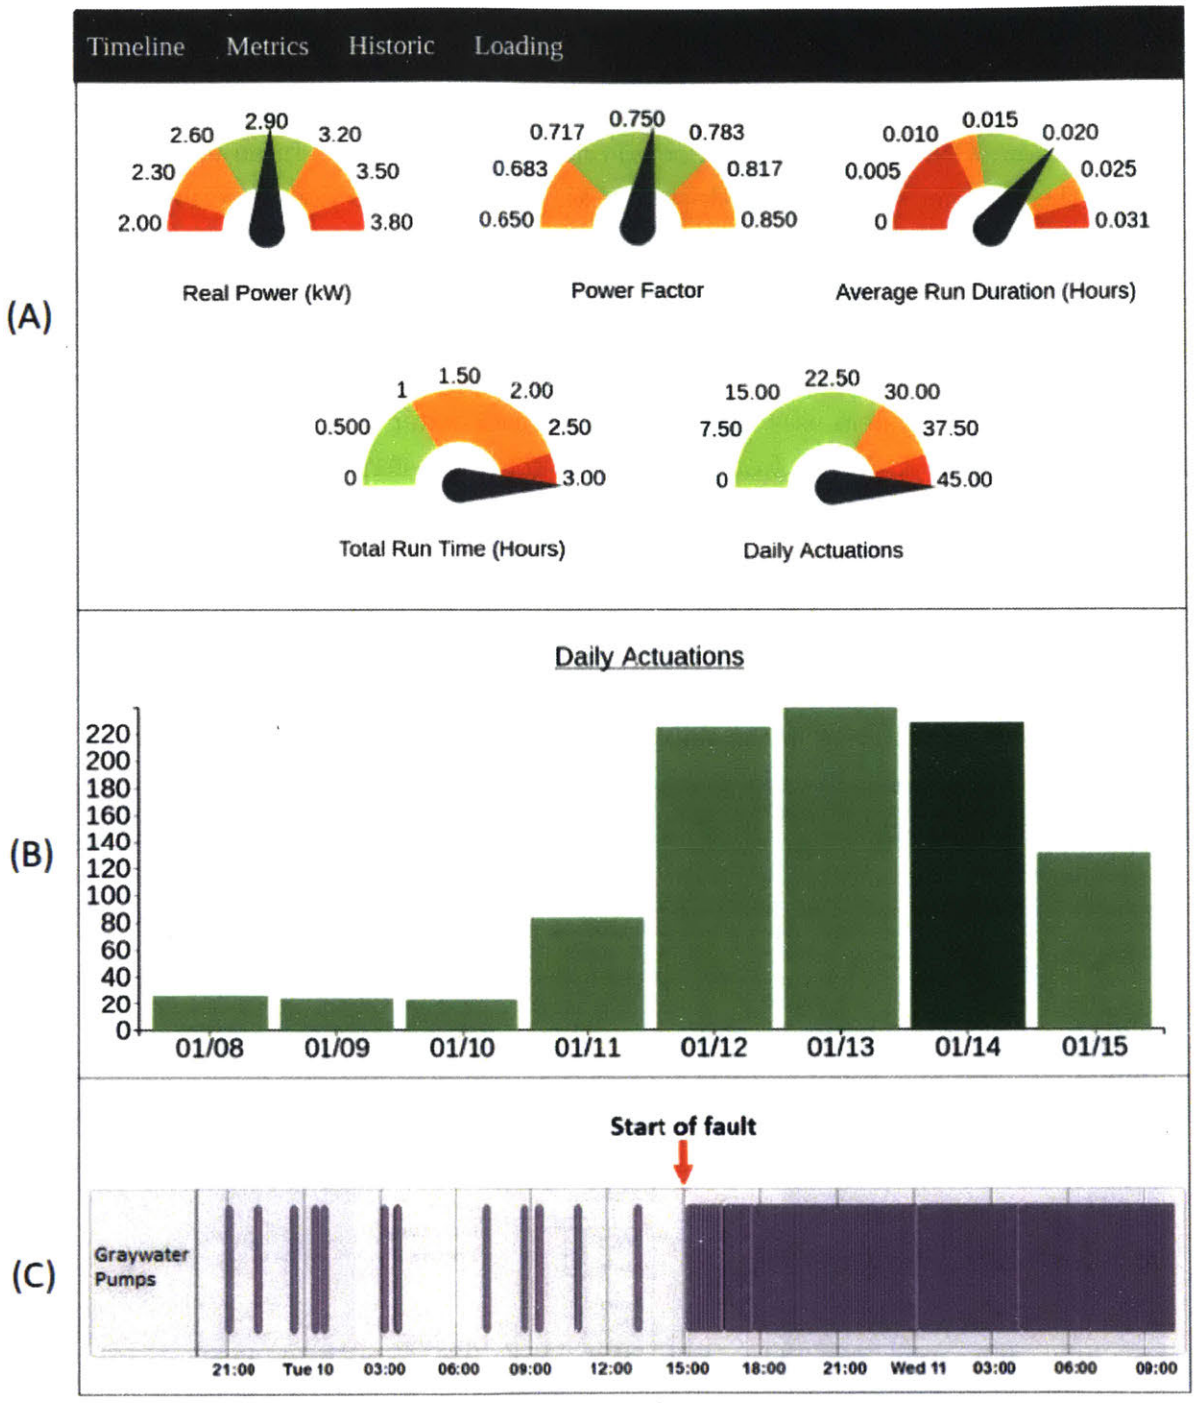

Figure 4-12: NILM Dashboard displays indicators of graywater sensor fault across three interfaces: (A) Metrics view (B) Historic view (C) Timeline view.

4.3 Sewage CHT System Faults

The installations aboard SPENCER and ESCANABA discussed until this point are just the latest iteration of NILM research conducted onboard WMEC-270s. In an effort to broadly examine the fault detection capabilities of NILM, the next two sections describe noteworthy faults detected in past studies [37–40].

The WMEC-270 collection, hold, and transfer (CHT) system processes sewage for a 100-plus person crew while underway, and a significant failure to this system can quickly force the cutter to shore for repairs. Some key elements of the CHT system are shown below in Figure 4-13. The collection system uses vacuum to draw black water from across the ship into the vacuum collection tank (VCT), from which the water is transferred to a larger holding tank. Two pumps maintain the vacuum in the VCT, monitored by pressure switches that sense vacuum through an orifice in the side of the tank. As sewage is collected, the vacuum slowly drops and the pumps turn on to maintain suction. The VCT also contains three electrodes that sense low level in the tank, high level, and a higher, alarm level. A controller operates a pair of discharge pumps (distinct from the vacuum pumps) that can empty the VCT to another collection tank lower in the ship.

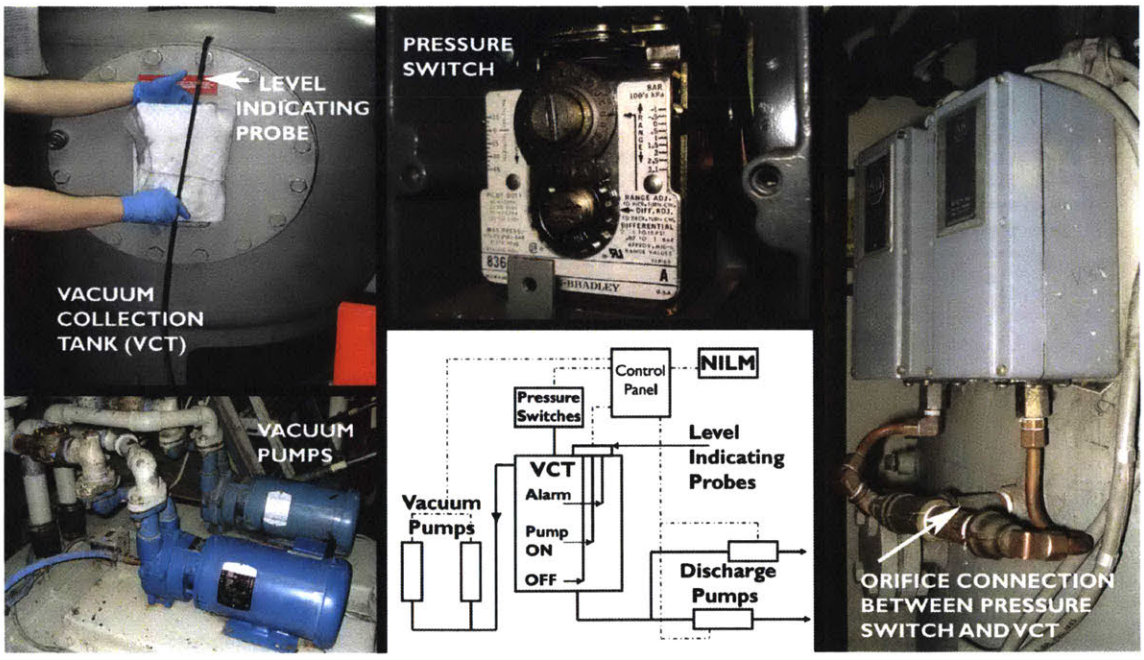

Figure 4-13: Collection, Hold, and Transfer (CHT) diagram and select photographs of components on USCGC SENECA.

Non-intrusive load monitoring was set up on CGC SENECA (WMEC-270) to observe the operation of the vacuum pumps, which had proven to be maintenance headache for the

crew. Normal CHT vacuum pump operation is show below in Figure 4-14. The vacuum pump actuates when pressure reaches 14inHg and secures when the pressure reaches 18inHg [40]. It should take a healthy vacuum pump about 60 seconds to restore vacuum to the desired level.

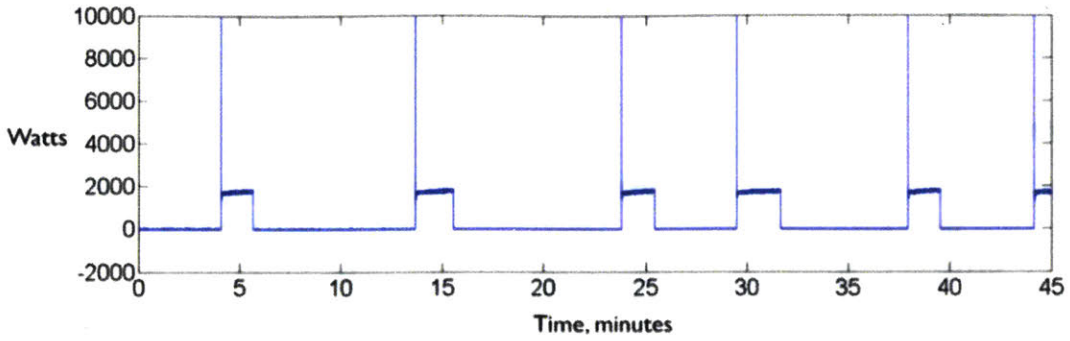

Figure 4-14: Healthy operations of the CHT vacuum pumps onboard CGC SENCECA [40]. Pump power demand is approximately 1800W and the pump runs for 60 seconds.

The highly corrosive nature of sewage and the numerous sensors required for proper operation make the system susceptible to soft faults. NILM observed the transient in Figure 4-15, which shows a vacuum pump run lasting roughly 20x the length a normal pump run. The NILM Dashboard has not been trained on CHT data, but we can theorize on its effectiveness as fault detection tool for this system. Like the graywater system fault described in the previous section, vacuum pump transient features in this fault condition remain identical to a healthy pump operation. The load identification and fault detection processes used in the NILM Dashboard system would detect this issue. The Run Duration and Total Run Time gauges on the Metrics would quickly enter the amber and red regions, alerting the watchstanders to an issue. Inspection of the system revealed that the extended pump runs were the result of a clogged vacuum pressure orifice, causing the pumps to run until vacuum reached nearly 24inHg.

Non-intrusive load monitoring on this system detected another suspect vacuum pump electrical transient, shown below in Figure 4-16. During this fault, the pump actuates due to low a vacuum reading, but is unable to draw a seal. After some time running without drawing the normal power or creating any vacuum, the back-up pump kicks in and restores vacuum to an acceptable level. This fault condition causes the pump transient to change shape, which makes could affect the ability of the NILM Dashboard to identify the fault and alert watchstanders. Depending on the other loads being monitored, the loss of pump seal may be incorrectly classified. However, with the electrical features of a fault condition identified, the ANN could be trained to classify a pump running with a very low real power levels as a

vacuum pump [36]. Another approach would be to pair the ANN with a correlation-based algorithm [41] to detect the specific shape of the double-pump run shown in Figure 4-16. Once the transient is properly identified, the low power levels and extended pump run times could be communicated to watchstanders via the Metric gauges.

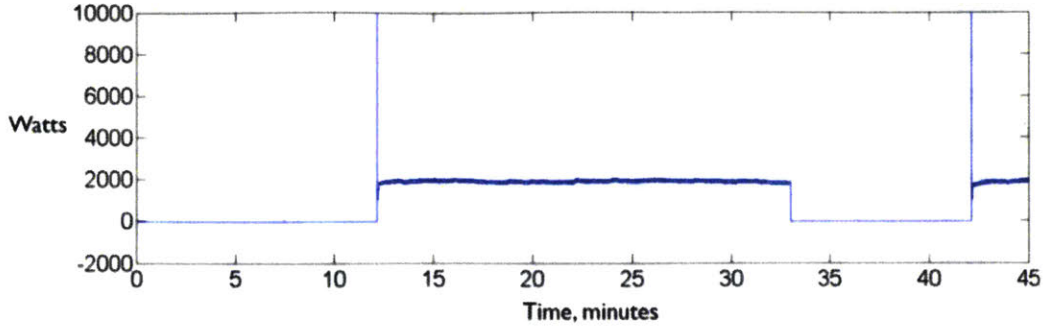

Figure 4-15: Extended CHT vacuum pumps run on SENCECA [37], caused by a clogged pressure switch orifice.

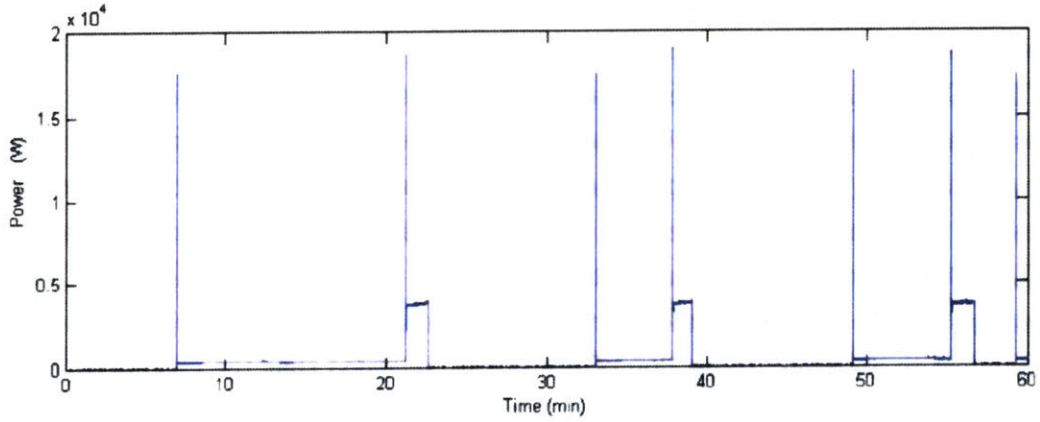

Figure 4-16: CHT vacuum pumps runs on SENCECA during a pump seal failure [40].

Both of these CHT faults are very difficult for even a dedicated watchstander to spot. The overall system continues to function, and a nominal examination of the CHT system in the fault condition shows no outward indications of a problem. However, the vacuum pumps run excessively, and can put the equivalent of 8 months of normal wear on the pumps in a matter of weeks [35]. When the soft fault eventually manifests as a hard fault because something has broken the system can become unavailable or, worse, flood an engineering space depending on the manifestation of the hard fault.

4.4 ASW Motor Coupling Fatigue

This section documents a fault detected on a WMEC-270 in a previous study, and is derived from analysis described in detail in reference [39]. Non-intrusive load monitoring was used to evaluate the mechanical coupling of the USCGC SENECA Auxiliary Sea Water (ASW) cooling pumps. When underway, saltwater cooling is provided to vital shipboard loads, such as the diesel generators and air conditioners, via two 40-hp, motor-driven centrifugal pumps. The coupling that joins the motor to the pump has a neoprene sleeve, designed to absorb shock and vibrations caused by misalignment and the motion of the ship. A picture of the ASW pump system and the associated coupling is shown below in Figure 4-17.

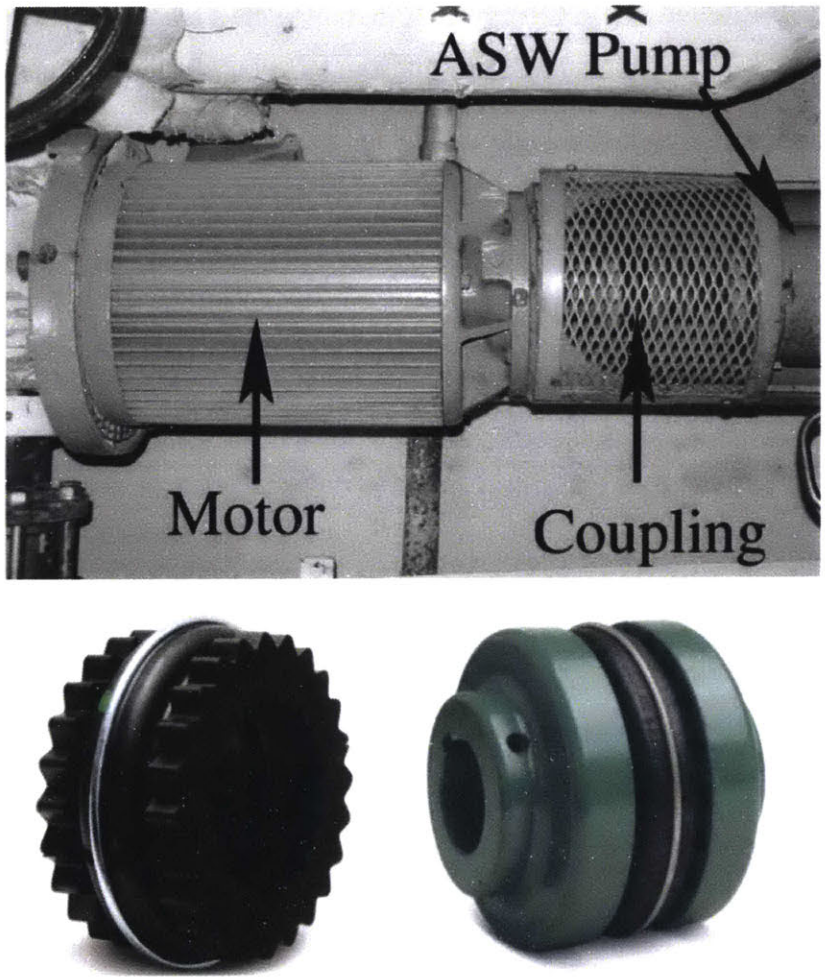

Figure 4-17: Top: WMEC Auxiliary Salt Water motor and pump system. Bottom: Flexible coupling system.

The ship's crew reported a series of coupling failures, leading to costly repairs and system

down time. A NILM sensor was placed upstream of the electrical panel servicing the ASW pumps. During the transient period when the pump activates, the power draw goes through a brief, high-frequency oscillation before reaching steady state. The coupling study focused on using these power oscillations as a diagnostic indicator. A physical model of the system approximates the coupling as a spring-damper system [39]. As the coupling degrades, the coefficients of stiffness and damping will change, causing this high frequency content to increase in magnitude prior to complete failure. To analyze this content, the following steps were used:

1. Compute the discrete Fourier transform (DFT) of the windowed segment of real power oscillations during motor start.
2. Scan DFT samples from 30-60Hz to find the location of the largest spectral peak.
3. Compute the energy in the band of frequencies centered around the spectral peak.

Figure 4-18 shows time-domain transient for a pump start, along with a DFT of the windowed data. A spectral peak is clearly present in the data at approximately 45 Hz. To evaluate the spectral peak as a diagnostic indicator, data was recorded on SENECA at five levels of increasing coupling deterioration. Figure 4-19 shows the spectral peak increasing in magnitude as the coupling wear increases.

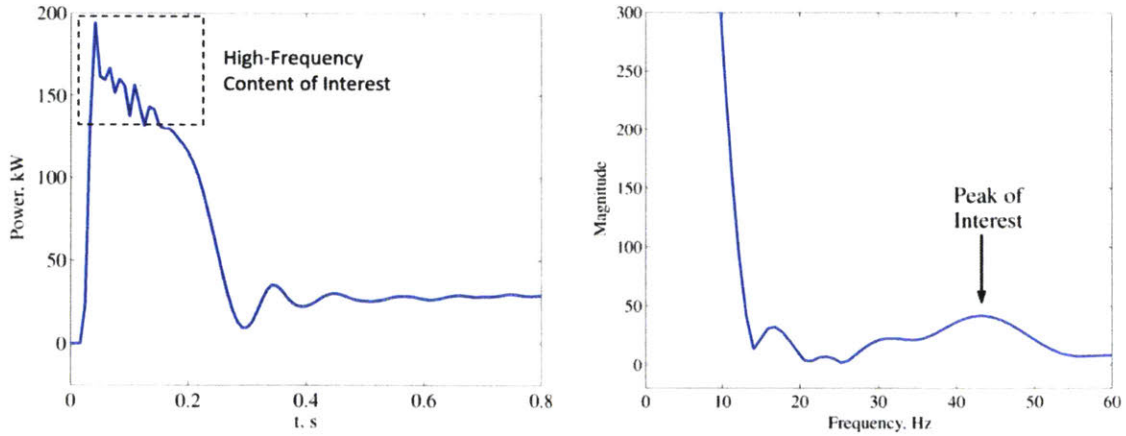

Figure 4-18: Left: Time domain view of motor power draw during typical start. A data vector is extracted from the high-frequency content and transformed. Right: DFT of motor start transient. Spectral peak is observed near 45 Hz [39].

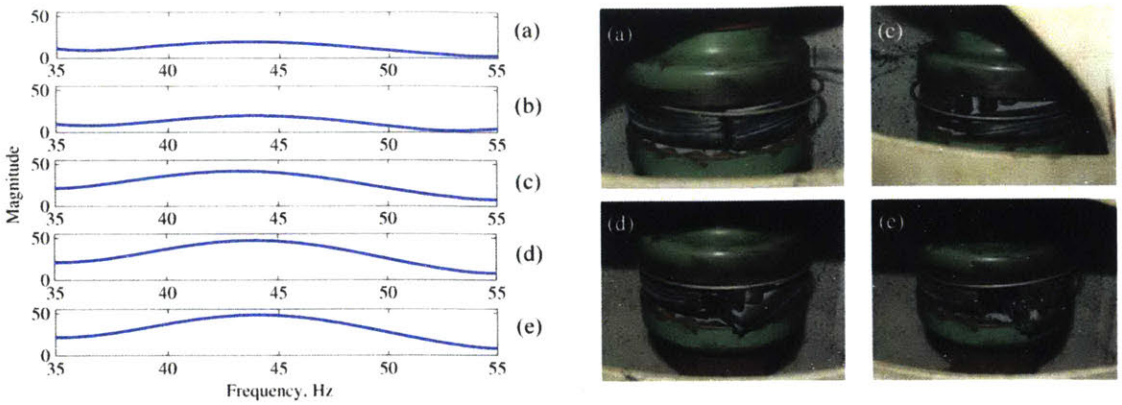

Figure 4-19: Left: Magnitude of frequency spectrum for motor starts at selected frequencies. Note the large jump between runs (b) and (c) Right: Coupling condition at each run [39].

The height of the spectral peak could serve as an additional Metric gauge on a future NILM Dashboard that monitors pumps with flexible couplings. As the spectral peak increases in magnitude, the gauge could alert watchstander that a replacement of the coupling may be necessary. Further analysis and additional case studies are needed before deploying this technique in the field and before widespread application to other coupling types. At present, a bench top experiment is underway to validate this approach for coupling fatigue assessment. This experiment will seek to examine spectral content of motor coupling systems across a wider variety of hardware types, wear profiles, and loading conditions.

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix A

Devices and Documentation

A.1 Non-Contact Calibrator

Figure A-1: Internal view of Non-contact Calibrator

Figure A-2: External view of NC Calibrator. Switch (far right) allows for power to be safely shifted between line-to-line or line-to-neutral phases.

System Schematic

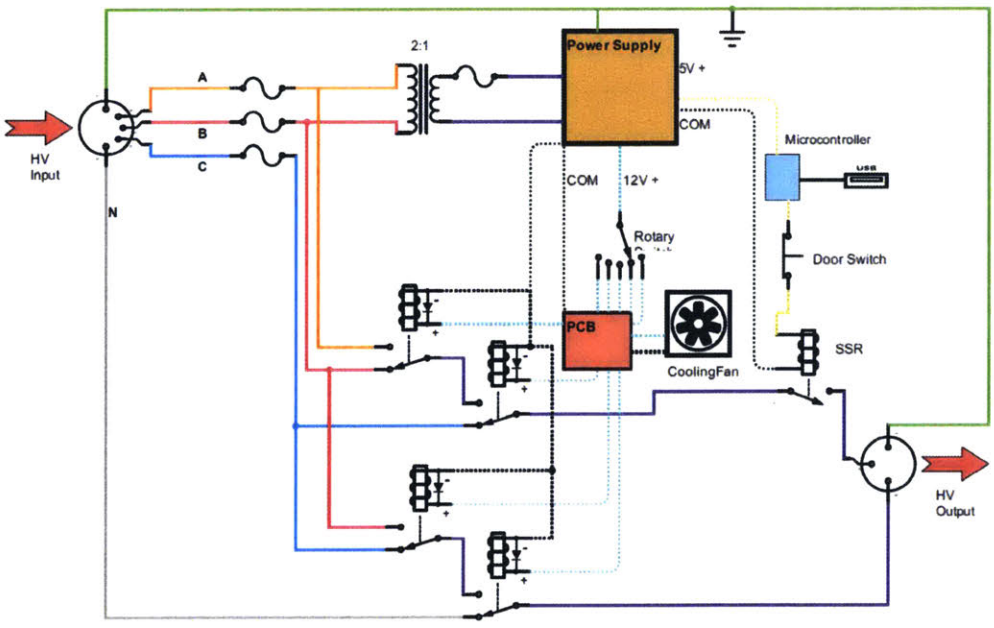

Figure A-3: The calibrator takes in 3-phase power and converts to a single phase for non-contact calibration.

PCB Diagram

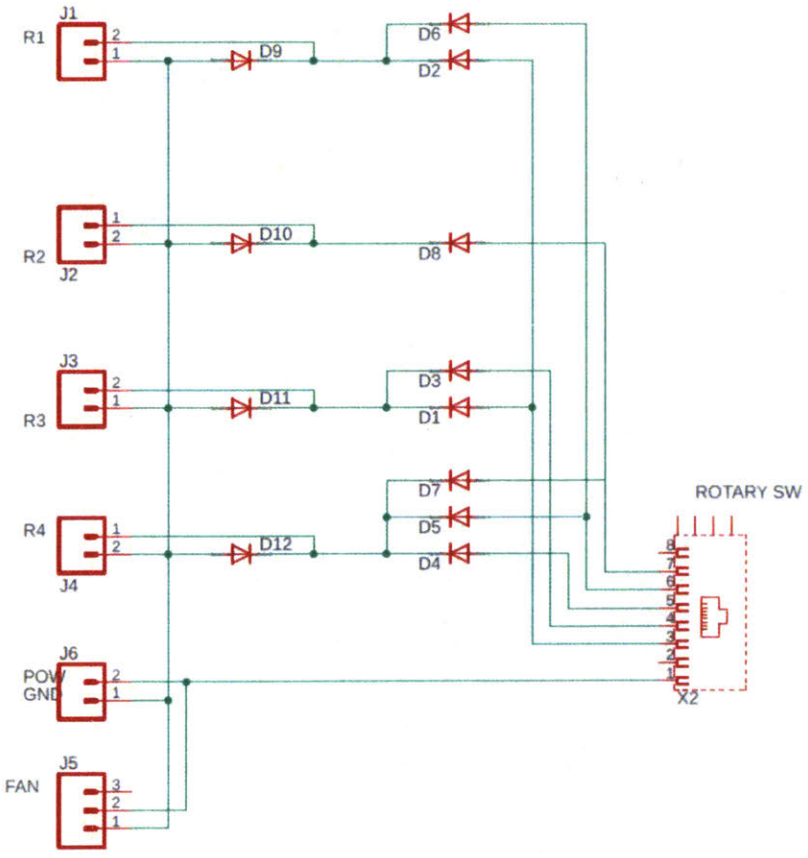

Figure A-4: PCB directs control voltage to the appropriate mechanical relay.

Calibration Load

Figure A-5: Load used for shipboard non-contact calibration, sonsists of two 2500W heaters.

Operation and Control

The non-contact calibrator safely converts a three-phase input current to a cyclic single-phase output current. The maximum safe voltage is 480 VAC. The 6-pole control knob allows the user to change from one phase to next during the calibration process. The clear panel on the top has a magnetic cut-out switch that allows for the safe changing of phases without interruption of the high voltage supply. The solid state relay in the calibrator is controlled by a standard WEMO smart plug PCB. The duty cycle of the output may be altered for different environments by connecting the calibrator to NILMbuntu computer. The standard NILM Plug CLI can be used to control the output cycles.

When onboard the ship, the calibrator is designed to be plugged into a welding outlet,

available in the DC shop or the forward passageway near the gun control room. In addition to the calibrator and 5000W load, there is a cord required to interface with the shipboard welding outlets. First, switch the shipboard outlet to the OFF position, indicated on the three-way switch below the outlet. Then plug the calibrator into the outlet, and then the heaters into the calibrator. Once the connection has been established, turn the outlet switch to ON position. The heaters should begin flashing. Check to ensure that the clear cover is fully closed so the magnetic switch is in the closed position. To switch between phases, simply open the cover and move the control knob.

Bill Of Materials

Part	Make	Model	Source	Cost	Quantity	Price
Power Supply	MEAN WELL	RD-50A	DigiKey	$ 20.50	1	$ 20.50
Phase Relay	OMRON	MGN1C-DC24	DigiKey	$ 35.42	4	$ 141.68
SSR Pattern Relay	Crydom	H12D4850	Digikey	$ 63.78	1	$ 63.78
Heat Sink	Crydom	HS122	DigiKey	$ 36.00	1	$ 36.00
Thermal Pad	Crydom	HSP1	Digikey	$ 1.20	1	$ 1.20
Recessed Cabinet	Arlington	DFVR3W-1	Amazon	$ 21.99	1	$ 21.99
Rotary Switch Knob	NKK Switches	HS16-3N-AT342	DigiKey	$ 21.74	1	$ 21.74
L2L Male Plug	McMaster	7162K68	McMaster	$ 22.97	1	$ 22.97
L2L Female Recept	McMaster	7162K86	McMaster	$ 39.62	1	$ 39.62
Power Recept	McMaster	7184K54	McMaster	$ 38.44	1	$ 38.44
Power Connector	McMaster	7184K42	McMaster	$ 61.44	1	$ 61.44
Ship Welding Plug	Todd	SYM 717.1 40A 450VAC 3P4	Todd Marine		1	$ -
Fan 80mm	Arctic	F8	Amazon	$ 7.99	1	$ 7.99
Fan Filter	SilverStone	FF81B	Amazon	$ 6.99	2	$ 13.98
Fuse Block	McMaster	7716K69	McMaster	$ 67.22	1	$ 67.22
Fuse 480V	McMaster	74485K892	McMaster	$ 9.36	4	$ 37.44
Enclosure	Integra	H141206S	Amazon	$ 126.99	1	$ 126.99
Aluminum Base	Integra	ABP1412	Amazon	$ 23.01	1	$ 23.01
Carrying Handle	Shoreline	170614	Amazon	$ 9.39	1	$ 9.39
Transformer	Hammond	PH100MQMJ	Allied	$ 61.94	1	$ 61.94
Transformer Fuse	McMaster	7714K12	McMaster	$ 10.15	1	$ 10.15
Magnetic Reed Switch	SoCell	B07PDV5K4C	Amazon	$ 8.58	1	$ 8.58
Disk Magnet	DIYMAG	HLMAG03	Amazon	$ 10.99	1	$ 10.99
Smart Plug PCB	**Remove PCB from WeMO plug, use for interface		LEES		1	$ -
Heater	Fostoria	MM-46A	Walmart	$ 193.36	1	$ 193.36
Heater Lamp	Fostoria	GF-3648H	Infrared Heate	$ 147.95	1	$ 147.95
Metal Guard	Fostoria	MMWG-46	Infrared Heate	$ 39.60	1	$ 39.60
						$1,227.95

A.2 NILM Dashboard Installation

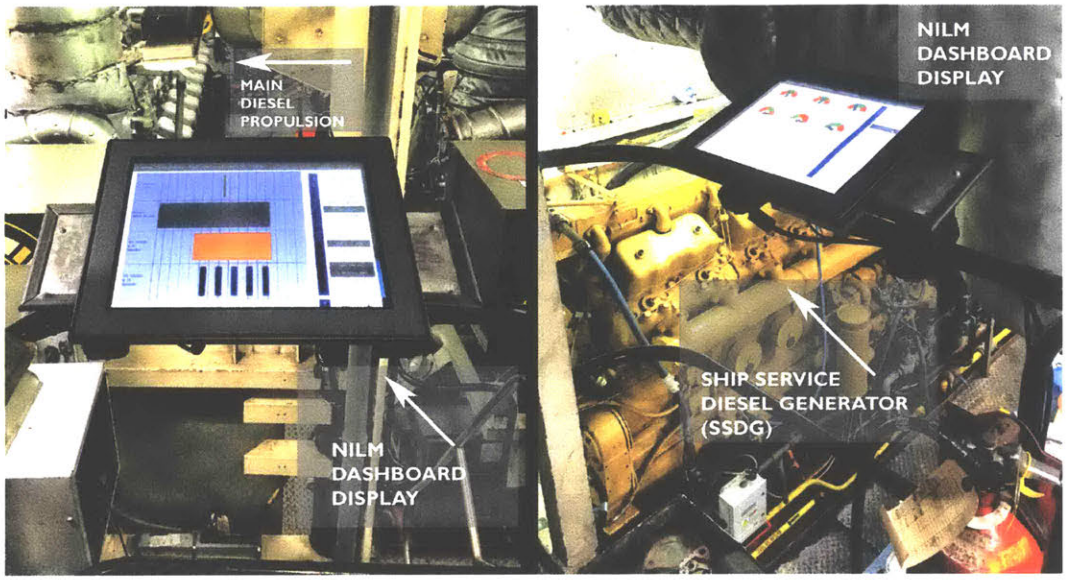

Figure A-6: NILM Dashboard display in SPENCER engine room.

Part	Make	Model	Vendor	Cost	Qty	Price
TouchScreen Monitor	NEXCOM	APPD1500T-WL VERB (10IAD1500003X0)	NEXCOM	$ 600.00	1	$ 600.00
Bracket Mount	ProSignal	B00BCRRUQ2	Amazon	$ 26.00	2	$ 52.00
Gasket	McMaster	4869A61	McMaster	$ 46.20	1	$ 46.20
Metal Plate	McMaster	89015K278	McMaster	$ 63.11	1	$ 63.11
						$ 761.31

A.3 Integrated NILM-Display (Kane-Box)

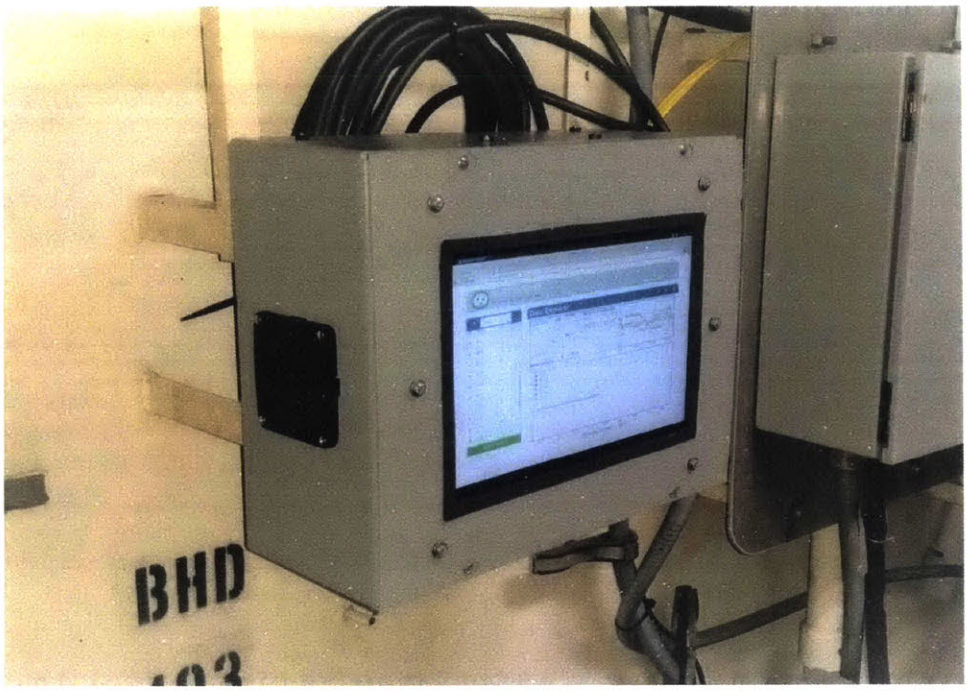

Part	Make	Model	Source	Cost	Quantity	Price
Metal Enclosure	Wiegmann	4DMV9	Grainger	49.38	1	$ 49.38
Surface Tablet (Refurb, Windows 10)	Microsoft	MQ2-00019-B	NewEgg	499.99	1	$ 499.99
**Must partition Surface HD for Ubuntu NILM						
External Hard Drive 1TB	WD Elements	WDBu6y0020bk-0A	Amazon	62.99	1	$ 62.99
USB Port (4)	Anker	AK-A7516011	Amaxon	9.99	1	$ 9.99
USB to Ethernet			GEM Lab	0	1	$ -
USB Fan	GDSTime	GDT8015S5V-USB	Amazon	12.98	1	$ 12.98
USB Panel Mount	CERRXIAN	B01FS4UB4M	Amazon	8.99	1	$ 8.99
Hard Drive Case	McMaster	69945K617	McMaster	27.28	1	$ 27.28
Power Supply Case	McMaster	69945K517	McMaster	22.28	1	$ 22.28
**Must Machine Cases for Cords						
Channel Gasket	McMaster	4869A76/4869A76	McMaster	33.8	1	$ 33.80
Contact NILM Box			GEM Lab	0	1	$ -
						$ 727.68

Manufacturing Notes:

This box is basically an enclosure that holds a contact NILM and a tablet running NILM-buntu. The enclosure from Grainger does not have a cut-out for the screen. It is strongly recommended to use MIT central machine to make the custom cutout. The hard drive/power supply cases from McMaster will work, but alternatives should be explored for Version 2.0. Also an internal hard drive is likely a better option if possible. Finally, the USB panel mount is on the bottom of the box in the current version but a top or side mount would be preferrable.

A.4 Magtrol Equipment

A.4.1 Basic Operation

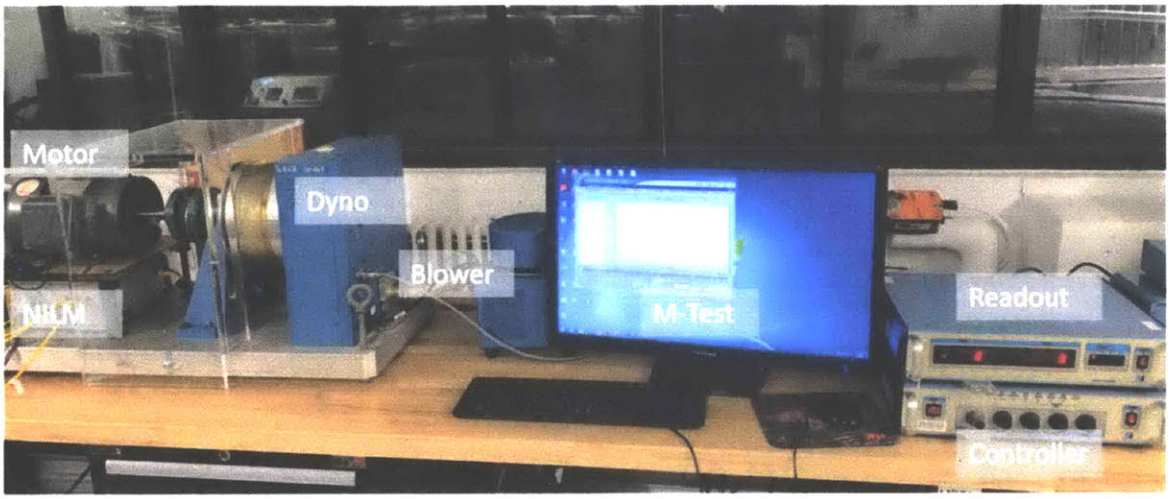

1. Download and install NI MAX Automation Explorer and Magtrol test software
2. Open NI MAX Automation & Instrumentation Explorer
3. From the left panel, select Devices and Interfaces > then NI-GPIB-USB-HS "GPIB0". This opens the window for the GPIB-USB cable
4. Click Scan for instruments and make sure it completes normally. If it doesn't restart controller and try again
5. Click on Interactive control, this will bring up a command window with GPIB0:
6. Type - IBFIND DEV9 – a prompt for DEV9 should appear - this writes to the controller primary address which is 9
7. From the DEV9 prompt write -*IBWRT*"M/r/n" - This should cause all the lights on the front panel of the controller to extinguish. This establishes remote control of the dyno.
8. If this fails, restart the controller and repeat steps 3-6
9. Open MTEST 5.0 from the desktop. Go to the Configure Hardware tab. The setup should load automatically but if it doesn't fill in the following fields:

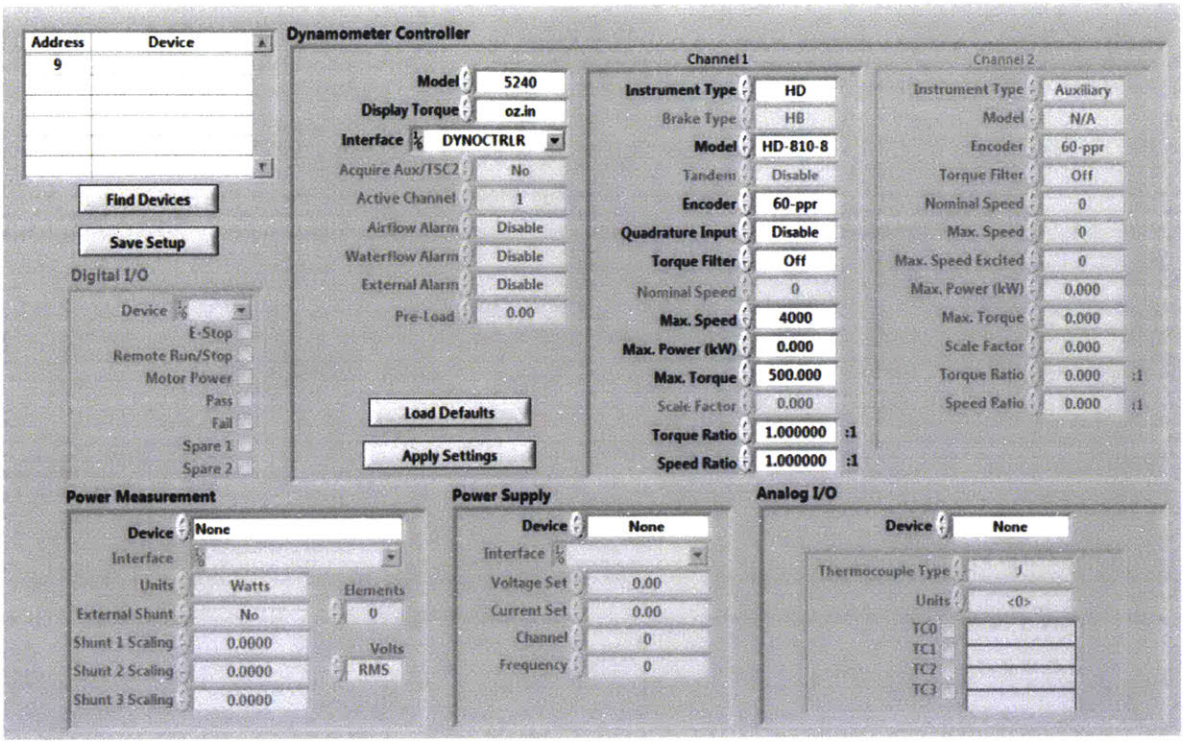

Figure A-7

10. Click "Save Setup". Go to the Display tab and make sure Speed, Torque, and Time are in the right window. Confirm any changes with "Save Setup".

11. Go to the Configure Test window. For a manual test, ensure the Control Parameter is Torque and the sample rate is .01 (seconds)

12. Go Test window. Before the motor runs, click Start Test. Make sure the window scrolls and the test is successful.

13. Plug in the motor and turn on the blower. Press start test. Turn on the motor. Adjust the torque with the "Load Control" bar.

14. Refer to the MTEST 5.0 manual and 5240 contoller manual for configuring other test types. Both are available on the Magtrol website.

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix B

NILM Dashboard

The NILM Dashboard provides a framework for a complete NILM system that captures data, accurately dissagregates load events, analyzes the equipment for potential faults, and presents useful information to end-users in real-time. This dashboard combines novel system architectures, algorithms, and user interfaces to solve not just the technical challenge of accurate load disaggregation, but also several technical challenges more commonly considered in the development of smart analytics [42] and industrial internet of things (IoT) devices [43]. Specifically, the NILM Dashboard simultaneously addresses the challenges of:

- Efficiently collecting, managing, and processing large volumes of raw electrical data
- Accurately disaggregating individual load operation
- Incorporating advanced algorithms for real-time monitoring and diagnostic focused analytics
- Providing users actionable insight into the operation of the electrical system and its individual loads through easy-to-understand visual displays

The NILM Dashboard is currently installed and operating onboard USCGC SPENCER. The system is located within the main enegine room and can easily be accessed by the crew, as shown in Figure B-1. Previous to this installation, fault detection and analysis has been conducted by MIT personnel after each cutter patrol. With the Dashboard in place, equipment analysis is brought to the cutter crew in real-time. This section describes the design and execution of the NILM Dashboard.

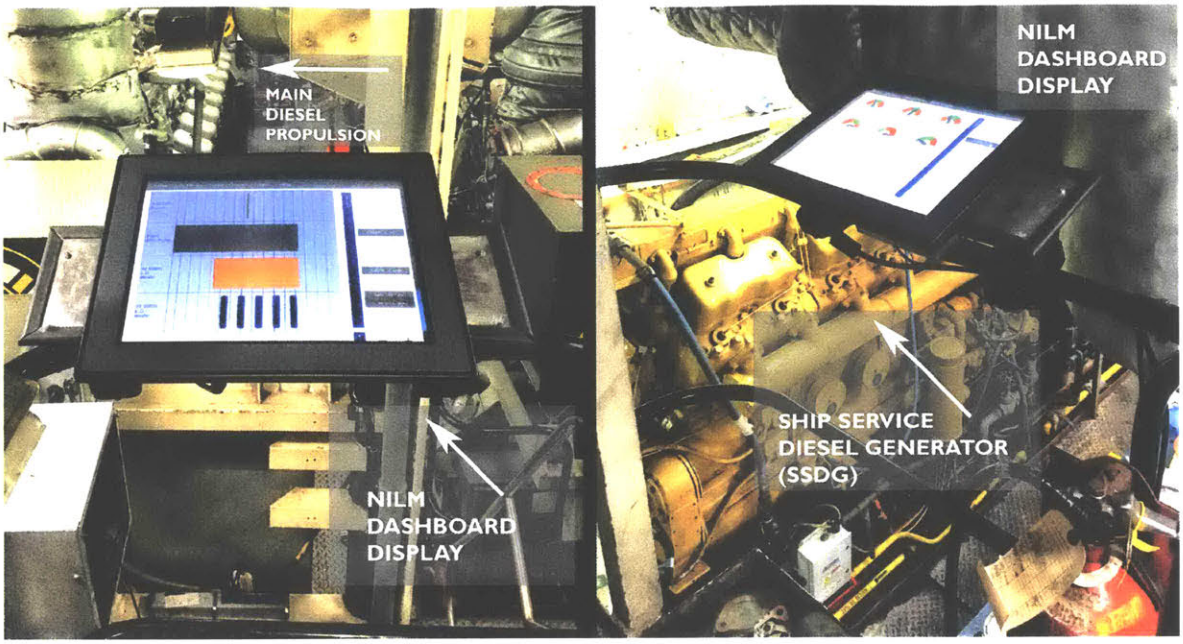

Figure B-1: Dashboard installation onboard SPENCER

B.1 System Architecture

The NILM Dashboard software stack runs on a typical Linux-based computer. The four layers shown in fig. B-2 are responsible for obtaining measurements, identifying loads, analyzing behaviors, and communicating results. This software stack is designed to address the technical challenges outlined in the introduction, including the ability to work with a large volume of data, accurately disaggregate electrical load events, provide real-time monitoring and diagnostic focused analytics, and the availability of an easy to understand visual display of information [42].

The Dashboard processes information about load *events*, which occur when equipment transitions between ON and OFF states. The event data is classified as a specific load and then mapped to an operating schedule of its activity. This time-based information enables detailed comparison of loads and provides a record of operation. In order to allow for predictive analytics, the platform also maintains load *metrics*, which are statistical conclusions that expose anomalies and patterns.

Data Capture and Pre-Processing

The NILM system captures currents and voltages from metering hardware at a sampling frequency of 3 kHz for non-contact sensors and 8 kHz for contact sensors. These high

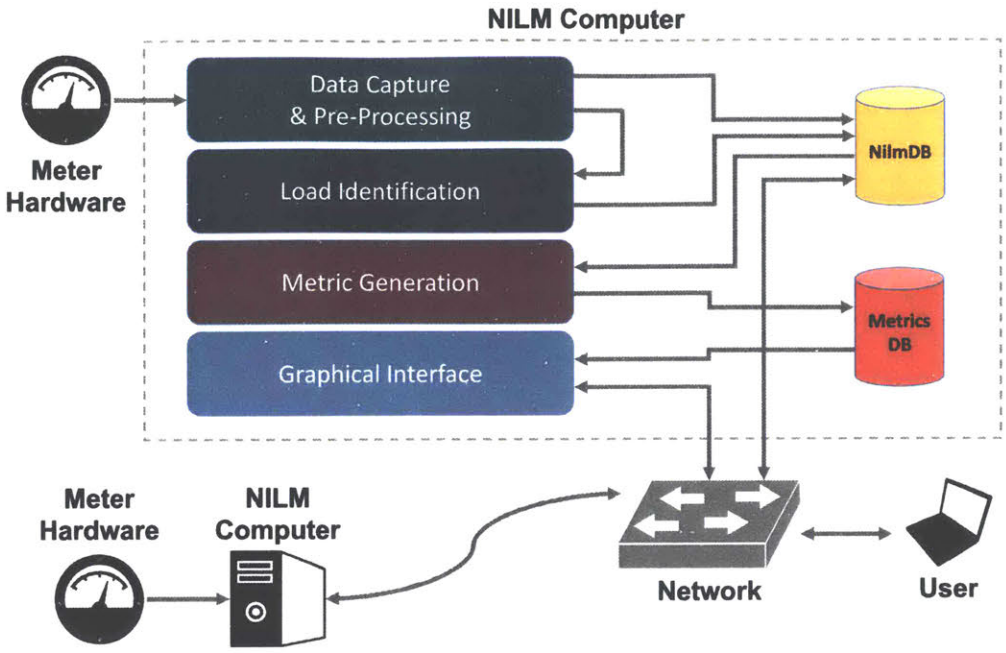

Figure B-2: System architecture of the NILM Dashboard. Rectangular blocks represent software layers, and cylinders represent storage databases. Lines between entities have arrows in the direction of data flow.

sampling rates are necessary for capturing transient shapes as loads change state [5] and the higher harmonic content of non-linear loads, e.g. variable frequency drives [6]. Sampling at these rates presents a data volume challenge in terms of processing, storage, and network bandwidth. To reduce this data volume while maintaining the relevant shape and harmonic information, these measurements are pre-processed into harmonic power envelopes using the Sinefit algorithm [7], effectively compressing the high-rate raw current and voltage data into real, reactive, and harmonic power components at a rate congruent with the power system line frequency (60 Hz for the ships). This promotes space-efficiency while maintaining the richness of the original signals. NilmDB is specially suited for storing this time-series data, making it available for high-speed and low-bandwidth access throughout the Dashboard platform [8].

Load Identification

The load identification block disaggregates the operation of individual loads from the power data in real-time. The NILM systems installed on the ships feature neural network architectures to achieve highly accurate load disaggregation, described in greater detail in Section

B.3. The resulting load events encode the type of state transition, as well as the change in real and reactive power levels. NilmDB records these sparse events at the times they occur. The Joule data processing framework [44] is used to streamline the capturing, pre-processing, and identification stages. This robust tool models the data pipeline as a series of processing "modules", with formally-defined "streams" of information passing between them. Data flows between modules through efficient memory pipes, without needing to access the database as an intermediary. This modular processing framework allows for efficient, real-time monitoring of load events.

Metrics Generation

To provide actionable diagnostics, the Dashboard must be able to assess the health of equipment in real-time and alert the user to anomalies. This crucial functionality makes the Dashboard an analytic tool, rather than simply a data collection device. To accomplish this, the metric generation block reads event streams from NilmDB and calculates operational metrics correlated to the health of systems and loads [14, 45]. Relevant metrics for a load may include the average power consumption or the number of operations per day. Calculated metrics and their associated metadata, such as type, load, and encompassed time range, are stored in MetricsDB for rapid query from the Dashboard. For the NILM systems on the Coast Guard cutters, six metrics are calculated. However, the MetricsDB must be able to adapt and expand for different NILM systems. The MetricsDB must also have the ability to save large data segments at full resolution, allowing for an anomalous load transient to be stored alongside its respective metrics. To meet these requirements, a NoSQL database structure is employed for MetricsDB [46, 47] due to its easy expandability and flexible size boundaries compared to traditional SQL alternatives.

Metric generation is implemented as a Python-based script that runs automatically every ten minutes. Every time the generator runs, it calculates *rolling* metrics over the past 24-hour window. These values are used to update the gauges in the metrics interface (Section III-B) to reflect current conditions. The metric generator also calculates *daily* load metrics if event data is available in NilmDB for the previous day, which ends at midnight. This information is used to populate the historic view (Section III-C) of the user interface, which shows the daily progression of a chosen metric for a given load.

Graphical Interface and Network

The final step is to transform load schedules and metrics into an interactive visual display that is user-friendly, easily-accessible and provides operators with actionable information.

This interface is described in detail in the next section. The Dashboard user interacts with a web-based application [30,31] that can be accessed on the NILM computer or from any computer connected on the same network. The user interface operates exclusively in the client's browser, conserving processing power and network bandwidth. The data visualizations are implemented with a data-driven graphical framework [48].

The browser must retrieve event and metric data to populate the interfaces. NilmDB has external web endpoints, so the client can directly communicate with this database to download event data for a particular load. The MetricsDB, however, cannot be accessed outside of the NILM computer. The web server proxies requests to the MetricsDB by exposing its own web endpoints for the client to reach. Many power systems are sufficiently complicated that more than one NILM system is required to track an entire electrical plant. The Dashboard software provides database access on a local network for NilmDB and MetricsDB, allowing a user to access all the NILM nodes on the local network from one place. The Dashboard is system is a self-contained, hard-wired network within the engine room of the ship. In its current form the Dashboard does not communicate or interface with any government systems.

B.2 User Interface

Providing the information generated by NILM quickly and clearly is paramount to creating a diagnostic tool capable of preventing failures to mission critical equipment. Even after processing, most NILM data is not intuitive to operators unaccustomed to analyzing equipment power streams and transients. The NILM Dashboard addresses this problem with an on-site interface that provides real-time and historic equipment information. To further aid operators in tactical decision making, the Dashboard generates useful metrics on system health. This information is made available to operators through an interface containing four interactive tools: Timeline, Metric View, Historic View, and Loading View.

Timeline

The "Timeline", as shown in Fig. B-3 provides a live view of the equipment status, allowing the user to see loads activated and secured in real-time. The user can monitor the entire plant or hide certain equipment from view, allowing for increased attention on select loads. The time window can be adjusted to the user's choosing, either through zoom/pan functions or by selecting one of four pre-set time periods. The Timeline tool provides the user with a compact picture of plant operations and the ability to easily investigate any apparent

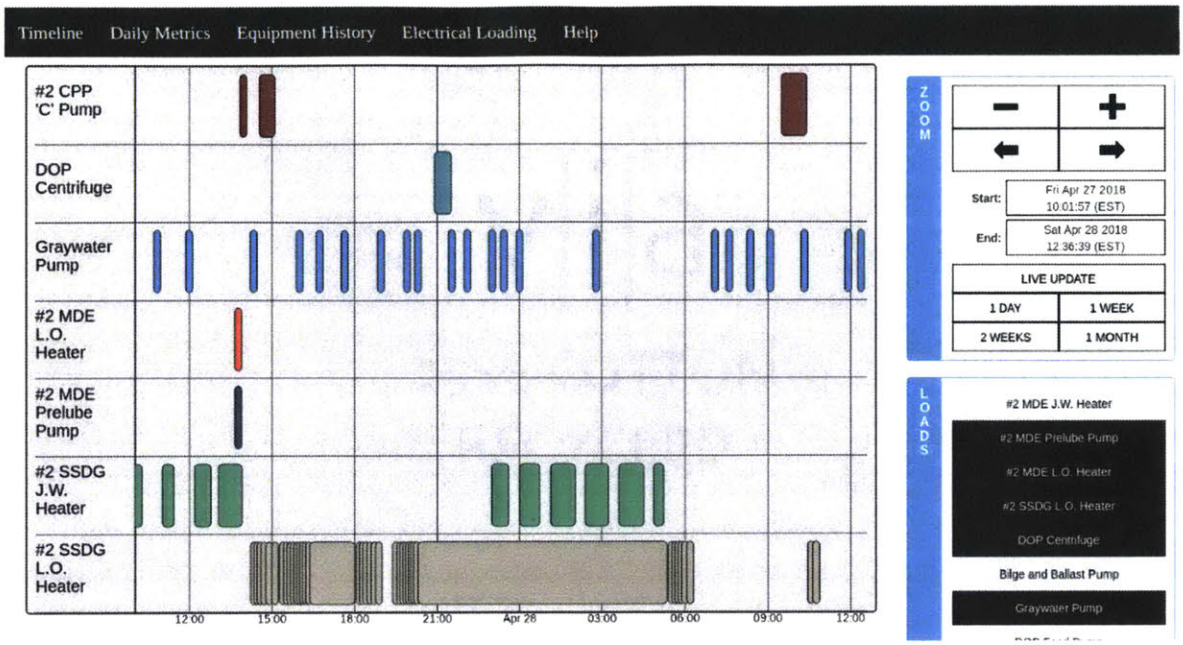

Figure B-3: Timeline interface displaying the status for a variety of engine room loads. Colored blocks represent periods when each load in energized.

anomalies.

Metric View

The "Metric View" shown in Fig. B-4 is a user's first stop when a fault is suspected. It provides the user with a set of diagnostic indicators for a selected piece of equipment. The metrics available are real power, apparent power, power factor, average run duration, total daily run time and daily number of actuations. Each metric is displayed as a gauge with green, yellow and red sections. The colored sections are derived from equipment nameplate data, known usage patterns and statistics from previous normal operation. Green indicates normal operations, while yellow and red indicate increasing likelihood of a fault. The gauge needle position is the average metric value for the last 24 hours and is refreshed every ten minutes. The Metric View provides an analysis of individual equipment health and helps direct initial troubleshooting efforts.

Historic View

The "Historic View" shown in Fig. B-5 provides short and long-term trend data to supplement the analysis from the Metric View. This tool allows users to select a single load and any one of the six metrics listed above. The Historic View is presented as a bar graph

Figure B-4: Gauges indicate whether the metrics pertaining to the selected load are within an acceptable range.

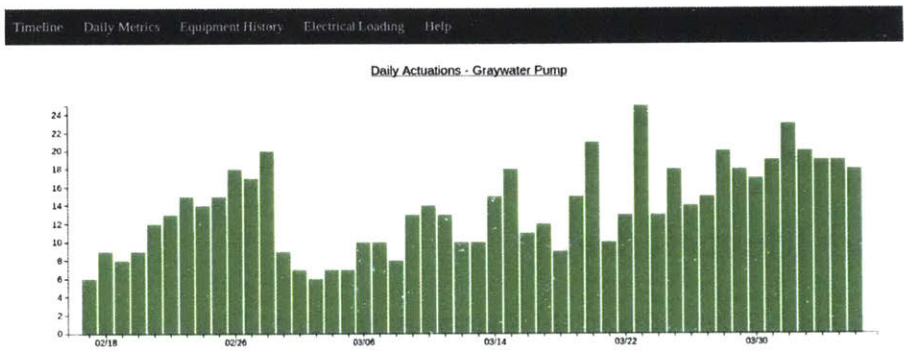

Figure B-5: The Historic View shows the daily trend for a given metric over a period of time.

that gives the user the ability to track equipment behavior over a period of up to 6 months. While the Metric View is intended for a watchstander to quickly detect a possible fault, the Historic View is designed for the plant manager to assess trend data, track behavior and make decisions on condition-based maintenance.

Loading View

The "Loading View" shown in Fig. B-6 allows a user to detect phase imbalances and loading discrepancies within the electrical system. The user can select a monitoring point and view the per phase power and total electrical load. In this case study, the Loading View allows the user to compare the total loading of the two generators or the two monitored subpanels. This information can be used for energy-scorekeeping and to optimize power generation.

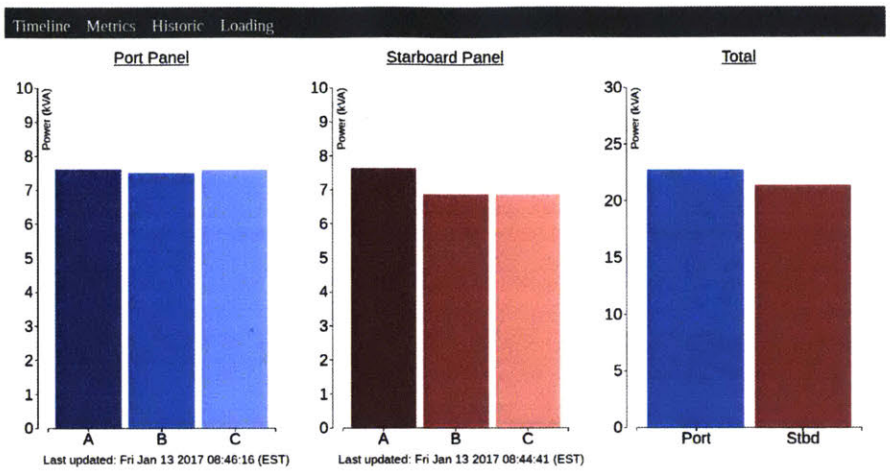

Figure B-6: The Loading View displays the current per-phase levels for a source of power data. With the shown configuration, the port and starboard subpanel are being compared.

B.3 Load Identification

Load identification is a key step in the four-stage pipeline. Load identification can be accomplished by many different algorithms, such as artificial neural networks (NN), k-nearest neighbors (k-NN), and correlation-based algorithms [8,9]. For this application, a neural network approach is taken. Since the two monitored subpanels have a fixed number of loads, a supervised learning approach is used in which data is hand-labeled in order to perform training. As described by Hart [15], there are three main categories of appliances that may be monitored by a NILM system: ON/OFF, Finite State Machine (FSM), and continuously variable. On the studied subpanels there are ON/OFF loads and one FSM load, the Diesel Oil Purifier. An ON/OFF load has only two states, ON or OFF, while a FSM load has several states due to its complex operation. Table 1.1 lists all the monitored loads which are to be displayed on the Dashboard.
Load identification occurs in the following three stages:

- **Event Detection** determines when transients occur in the power stream.
- **Event Classification** matches identity of transient to a load.
- **Load Confirmation** checks constraints between load events.

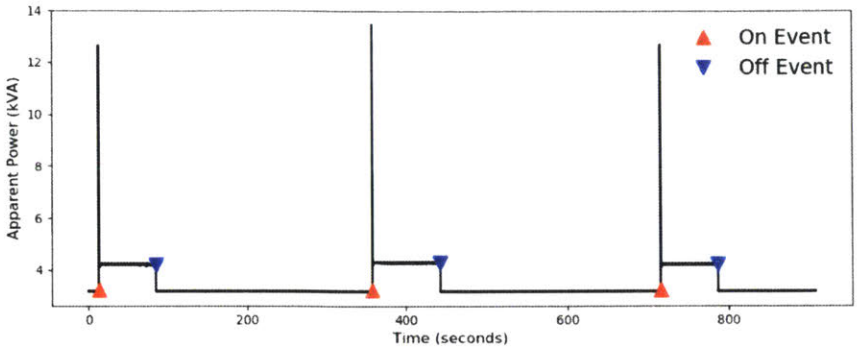

Figure B-7: ON-events and OFF-events for graywater pump runs.

Event Detection

After pre-processing, Real (P) and reactive (Q) power are outputted at every line cycle (60 Hz) for each phase. From the fundamental P and Q streams, apparent power (S) is calculated as,

$$S = \sqrt{P^2 + Q^2} \tag{B.1}$$

The S stream is then smoothed using a 101-point median filter, which eliminates small fluctuations while preserving edges. The S stream is used to detect turn-on and turn-off times by detecting where the stream abruptly changes in value. Converting to apparent power simplifies load detection to a single data stream. The filtered data is convolved against the Laplacian of a Gaussian [49] kernel to compute the smoothed second-derivative of this stream. This effectively maps step changes in apparent power to zero-crossings for easier detection. An empirically-determined threshold is set to remove zero-crossings that are due to small fluctuations of the resulting convolution. A zero-crossing detector is then used to find the location of the steps. Fig. B-7 shows several ON and OFF events as detected by the edge detector.

After an ON or OFF event is detected using the S stream, the edge detector examines the P and Q to calculate a set of features to be used as an input vector to the NN. For each phase (A,B,C), represented here as x, an ON event produces four features, the transient-peak ($P_{on,x,peak}$, $Q_{on,x,peak}$) and the steady-state level changes ($P_{on,x,ss}$, $Q_{on,x,ss}$). Because there is no transient peak when a load turns off, an OFF event has only two features per phase ($P_{off,x,ss}$, $Q_{off,x,ss}$). The steady state level change is calculated by taking the difference in the median power level for a defined window (Δt_{window}) before and after an ON/OFF event. The Δt_{window} is chosen to be 0.5 sec, which is 30 samples given a line frequency of 60 Hz. Fig. B-8 illustrates these features for the real power stream on one phase. Features for reactive power are calculated using the same process. Equations (B.2)-(B.6) generate the features

and the input vectors.

ON event steady state change in power:

$$\begin{aligned} P_{on,x,ss} = median(P_x[t_{end}], ..., P_x[t_{end} + \Delta t_{window}]) \\ -median(P_x[t_{on} - \Delta t_{window}], ..., P_x[t_{on}]) \end{aligned} \tag{B.2}$$

ON event transient peak:

$$\begin{aligned} P_{on,x,peak} = max(P_x[t_{on}], ..., P_x[t_{end}]) \\ -median(P_x[t_{end}], ..., P_x[t_{end} + \Delta t_{window}]) \end{aligned} \tag{B.3}$$

OFF event steady state change in power:

$$\begin{aligned} P_{off,x,ss} = median(P_x[t_{off}], ..., P_x[t_{off} + \Delta t_{window}]) \\ -median(P_x[t_{off} - \Delta t_{window}], ..., P_x[t_{off}]) \end{aligned} \tag{B.4}$$

Here, t_{on} is the time the load turns on, t_{end} is the end of the start-up transient, t_{off} is the time the load turns off, and t_{window} is the length of the window for taking steady state calculations.
The feature input vector for ON events is:

$$\begin{aligned} (&P_{on,A,ss}\ Q_{on,A,ss}\ P_{on,A,peak}\ Q_{on,A,peak} \\ &P_{on,B,ss}\ Q_{on,B,ss}\ P_{on,B,peak}\ Q_{on,B,peak} \\ &P_{on,C,ss}\ Q_{on,C,ss}\ Q_{on,C,peak}\ Q_{on,C,peak}) \end{aligned} \tag{B.5}$$

The feature input vector for OFF events is:

$$\begin{aligned} (&P_{off,A,ss}\ Q_{off,A,ss} \\ &P_{off,B,ss}\ Q_{off,B,ss} \\ &P_{off,C,ss}\ Q_{off,C,ss}) \end{aligned} \tag{B.6}$$

There are a fixed number of loads on the panels and the normal real and reactive power draw of each load is known. Therefore the total change in steady state should not be less than the smallest load on the panels. False event detections are reduced by comparing the magnitude of the calculated change in steady state values to a threshold. Events with power value changes beneath this threshold are discarded. In one month (September 2017), there were a total of 2946 events on the USCGC Spencer port panel. The described edge detector

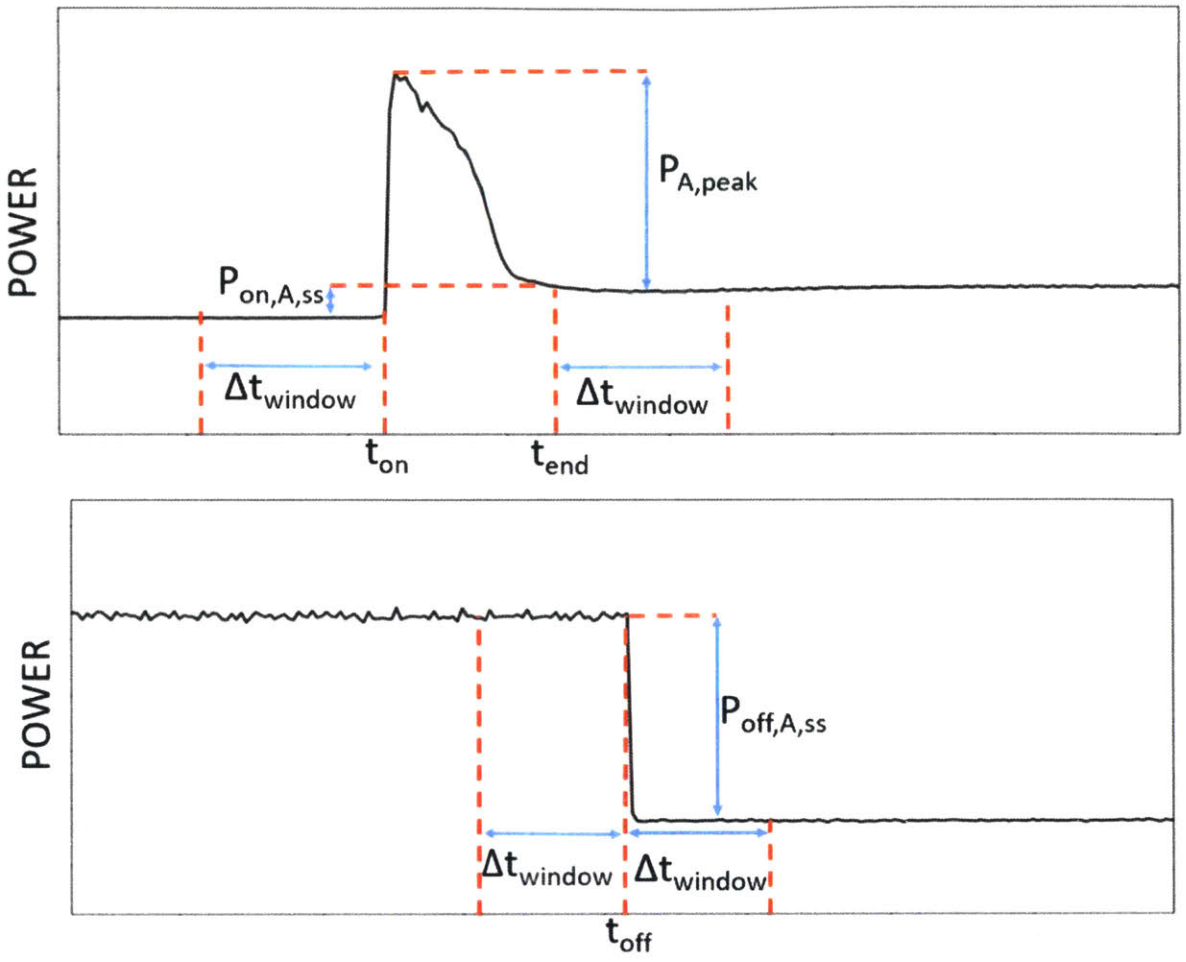

Figure B-8: ON/OFF load features for a single phase

correctly found 2936 of those events and incorrectly detected 32 non-events.

Event Classification

To classify each event as an individual load, a fully-connected neural network (NN) using stochastic gradient descent (SGD) is applied [36]. Separate NNs are utilized for ON events and OFF events for each panel. The input layer to the NNs are the features for each load as previously described in (B.5) for ON events and (B.6) for OFF events. Two hidden layers of 10 nodes each are used, for which a weighted sum of the inputs from the previous layer to each node is passed through a ReLu (Rectified Linear Unit) [50] activation function, $f(z) = max(0, z)$, where z is the input to a node. Back-propagation is used to find the set of weights and biases to minimize loss. To allow multi-class classification, the output layer is a softmax layer, [51], defined as,

$$f(z)_i = \frac{e^{z_i}}{\sum_{j=1}^{k} e_j^z} \text{ for } i = 1...k \tag{B.7}$$

where z is the input to each node, and k is the number of classes. Each ON/OFF load is a unique class. Additionally, some loads often actuate together during normal operations, as they are part of a combined system, creating a new class representing multiple loads. Finally, the Diesel Oil Purifier has a class for each state of operation.The output is a vector of probabilities that sums to one and the classification is made by selection of the class that has the highest probability.

The data was split into three sets: training, validation, and testing. The prediction error of the validation set was used as a stopping criterion during training [52]. To prevent over-fitting, a third separate data set, the testing set was used. Repeated random sub-sampling was performed on the data for 10 iterations, so that for each iteration the data was randomly split by load class into training, validation, and testing data and the performance was evaluated on the testing data.

Some loads have short runs, in the range of minutes, and thus have many events. Other loads have run durations in the range of hours or even days, and thus have far fewer events. To prevent the NN from simply predicting the most common class, the training data had to be better balanced. This was accomplished by random under-sampling of the majority classes on the training and validation data.

It is not sufficient to focus only on the total percentage of correctly classified loads when verifying the accuracy of the identifier. If a load only turns on and off a few times over a month, incorrectly identifying it will not have much effect on the total classification accuracy, but it is still vital that the load be correctly identified. Thus, the accuracy of the model for each class is evaluated by considering the following parameters [53]:

- True Positive (TP): a load event occurs and is correctly identified
- False Positive (FP): a load event is classified, but that event did not occur
- False Negative (FN): a load event occurs but that event is not classified

These parameters are used to determine the classifier's recall and precision, which answer two fundamental questions:
(1) What is the likelihood that a load event is reported? (*recall*)
(2) What is the likelihood that a reported event is correct? (*precision*)

$$recall = \frac{TP}{TP + FN} \tag{B.8}$$

Table B.1: Accuracy of classifying on-events

Load	**TP** $\pm \sigma_{TP}$	**precision** $\pm \sigma_P$	**recall** $\pm \sigma_R$
Main diesel engine (MDE) keep-warm system			
LO Heater	8 ± 0	0.92 ± 0.07	1 ± 0
JW Heater	1 ± 0	0.95 ± 0.16	1 ± 0
Prelube Pump	4 ± 0	0.80 ± 0.19	1 ± 0
Ship service diesel generator (SSDG) keep-warm system			
JW Heater	65 ± 0	0.97 ± 0.02	1 ± 0
LO Heater	95 ± 0	0.98 ± 0.01	0.99 ± 0.01
Additional engine room loads			
CPP Pump	4.2 ± 1.5	0.88 ± 0.14	0.84 ± 0.15
Graywater Pump	268.3 ± 1.06	0.99 ± 0	1 ± 0
DO Purifier (Centrifugal)	26 ± 1.05	0.98 ± 0.04	1 ± 0
DO Purifier (Feed Pump)	7 ± 0	1 ± 0	1 ± 0
DO Purifier (Flushing Sequence)	13.9 ± 0.32	0.95 ± 0.04	1 ± 0

$$precision = \frac{TP}{TP + FP} \tag{B.9}$$

Table B.1 shows the results for each load from the 10 iterations. The results presented are the average number of true positives, the precision, and the recall for each load, as well as the standard deviations of each (σ_{TP}, σ_P, and σ_R respectively). Precision and recall values of one with a standard deviation of zero indicate perfect performance in identifying a specific class. For true positives, a small standard deviation shows that performance between iterations is consistent and that the model is not overfitting.

Confirmation and Implementation

Load identification is implemented into the NILM Dashboard architecture through two modules, classification and confirmation. The classification module runs the previously described detection and classification algorithm on 10 second windowed power stream data. Since there may be events that occur at the edges of an interval, the intervals are overlapped to ensure that no events are lost. For each iteration, the classification module output is piped to a confirmation module. For each load, the confirmation module stores the most recent state (ON or OFF) outputted to NilmDB, to ensure two ON events or two OFF events are not output consecutively. If two consecutive ONs are detected, the confirmation module removes the first occurrence from the NilmDB. If two consecutive OFFs are detected, the second oc-

currence is not outputted to the NilmDB. This reduces the possibility of the Dashboard incorrectly displaying that a load is energized.

B.4 Observations and Conclusion

For greater classification accuracy, there is further research being done on integrating this NN method with other classification methods such as an exemplar shape-matching algorithm [8] and multi-scale median filtering. This would all be done within the classification module, allowing the classification algorithm to be updated without affecting the operation of the user interface.

The next iteration of the NILM Dashboard will use the known loads to directly infer the status of the propulsion plant, or even the whole ship. For instance, the jacket water heaters, lube oil heaters and pre-lube pumps can be used to determine whether the main diesel engine is online, secured or in standby. The main diesel engine essentially becomes a finite state machine, with the stages of operations determined by NILM monitored equipment. With the right indicators, even the operational status of the entire ship can be determined from NILM. For example, the controllable pitch propeller (CPP) pumps are energized when the ship enters a higher state of readiness known as the restricted maneuvering doctrine (RMD). The Timeline View can then display when the ship enters RMD.

The Dashboard platform can be adapted to any NILM system, whether it be in a house, factory, or naval vessel, to provide feedback on equipment behavior and energy usage. The NILM Dashboard provides the framework and analysis tools to turn power stream data into actionable information for optimizing operations.

Bibliography

[1] G. Bredariol, "The shipboard automatic watchstander (saw): Utilization of nonintrusive load monitoring for shipboard automation," Master's thesis, Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, MA, 6 2017.

[2] W. Cotta, "Machinery diagnostics and characterization through electrical sensing ," Master's thesis, Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, MA, 6 2015.

[3] J. S. Donnal, J. Paris, and S. B. Leeb, "Energy applications for an energy box," *IEEE Internet of Things Journal*, vol. 3, no. 5, pp. 787–795, 2016.

[4] D. Green, "Framework of non-intrusive load monitoring for shipboard environments," Master's thesis, Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, MA, 6 2018.

[5] 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.

[6] 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.

[7] 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.

[8] J. Paris, J. S. Donnal, and S. B. Leeb, "NilmDB: The non-intrusive load monitor database," *IEEE Transactions on Smart Grid*, vol. 5, no. 5, pp. 2459–2467, 2014.

[9] A. Zoha, A. Gluhak, M. A. Imran, and S. Rajasegarar, "Non-intrusive load monitoring approaches for disaggregated energy sensing: A survey," in *Sensors*, 2012.

[10] 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.

[11] A. Coraddu, L. Oneto, A. Ghio, S. Savio, D. Anguita, and M. Figari, "Machine learning approaches for improving condition-based maintenance of naval propulsion plants," *Engineering for the Martime Environment*, vol. 230, pp. 136–153, 2016.

[12] V. V. Sourabh Dash, "Challenges in the industrial applications of fault diagnostic systems," *Computers and Chemical Engineering*, vol. 24, pp. 785–791, 2000.

[13] R. Beebe, "Estimate the increased power consmption caused by pump wear," *Pump Magazine*, vol. 58, pp. 20–27, 2008.

[14] J. C. Nation, A. Aboulian, D. Green, P. Lindahl, J. Donnal, S. B. Leeb, G. Bredariol, and K. Stevens, "Nonintrusive monitoring for shipboard fault detection," in *Sensors Applications Symposium (SAS), 2017 IEEE.* IEEE, 2017, pp. 1–5.

[15] G. W. Hart, "Nonintrusive appliance load monitoring," *Proceedings of the IEEE*, vol. 80, no. 12, pp. 1870–1891, 1992.

[16] W. A. Shewhart. American Society for Quality (ASQ), 1980. [Online]. Available: https://app.knovel.com/hotlink/toc/id:kpECQMP002/economic-control-quality/economic-control-quality

[17] S. Rao, *Reliability-Based Design.* Mcgraw-Hill, 1992.

[18] W. A. Levinson, *Statistical process control for real-world applications.* CRC Press, 2011.

[19] Technical Manual 2813-241-A: Lube Oil System, accessed: 2019.

[20] Surface Technical Information Portal, https://www.uscg.mil/, accessed: Jan 2019.

[21] H. Woud and D. Stapersma, *Design of propulsion and electric power generation systems.* Institute of Marine Engineering, Science, and Technology, 2002.

[22] Technical Manual 2819-302-B: Heater Contollers, accessed: 2019.

[23] Technical Manual 2814-233-A: Main Diesel Engine, accessed: 2019.

[24] Technical Manual 2824-593-B: Waste Water System, accessed: 2019.

[25] Technical Manual 2817-245-A: Controllable Pitch Propellor System, accessed: 2019.

[26] Technical Manual 2833-541-A: Oil Purifier System, accessed: 2019.

[27] Technical Drawing 905-WMEC 529-001: Ballasting and Emergency Bilge Drainage System, accessed: 2019.

[28] R. Faturechi, M. Rose, and C. Miller, "Years of warnings, then death and disaster: How the navy failed its sailors," *ProPublica*, Feb 2019. [Online]. Available: https://features.propublica.org/navy-accidents/us-navy-crashes-japan-cause-mccain/

[29] "Future proofed: Autonomous shipping research," *Nautilus International*, Feb 2018. [Online]. Available: https://www.nautilusint.org/en/news-insight/resources/nautilus-reports/autonomous-shipping-research/

[30] "Node.js," https://nodejs.org/, accessed: Jan 2018.

[31] "Express – node.js web application web framework," https://expressjs.com/, accessed: Jan 2018.

[32] "Comdtinst m3123.13: Operational reporting," 2014.

[33] A. Cowan, "Review of recent commerical rooftop unit field studies in the pacific northwest and california," Tech. Rep., Oct 2004.

[34] M. Breuker and J. Braun, "Common faults and their impacts on rooftop air conditioners," *International Journal of HVAC + R Research*, vol. 4, pp. 303–318, June 1998.

[35] S. Leeb, P. Lindahl, D. Green, T. Kane, J. Donnal, and S. Kidwell, "Power as predictor and protector," *Marine Technology*, pp. 28–35, April 2019.

[36] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 05 2015. [Online]. Available: http://dx.doi.org/10.1038/nature14539

[37] R. W. Cox, M. Piper, G. Mitchell, P. Bennett, J. Paris, W. Wichakool, and S. Leeb, "Improving shipboard maintenance practices using non-intrusive load monitoring," in *ASNE Intelligent Ship's Symposium*, 2007.

[38] R. W. Cox, "Minimally intrusive strategies for fault detection and energy monitoring," Ph.D. dissertation, Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, MA, 2006.

[39] S. R. Shaw, S. B. Leeb, L. K. Norford, and R. W. Cox, "Nonintrusive load monitoring and diagnostics in power systems," *IEEE Transactions on Instrumentation and Measurement*, vol. 57, no. 7, pp. 1445–1454, July 2008.

[40] M. Piber, "Improving shipboard maintenance practices using non-intrusive load monitoring," Master's thesis, Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, MA, 6 2007.

[41] S. Leeb, S. Shaw, and J. Kirtley, "Transient event detection in spectral envelope estimates for nonintrusive load monitoring," *Transactions on Instrumentation and Measurement*, vol. 10, pp. 1200–1210, 1995.

[42] D. Alahakoon and X. Yu, "Smart electricity meter data intelligence for future energy systems: A survey," *IEEE Transactions on Industrial Informatics*, vol. 12, no. 1, pp. 425–436, Feb 2016.

[43] L. D. Xu, W. He, and S. Li, "Internet of things in industries: A survey," *IEEE Transactions on Industrial Informatics*, vol. 10, no. 4, pp. 2233–2243, Nov 2014.

[44] J. Donnal, "Joule: Decentralized data processing," http://docs.wattsworth.net/joule/, accessed: Jan 2018.

[45] J. Paris, J. S. Donnal, R. Cox, and S. Leeb, "Hunting cyclic energy wasters," *IEEE Transactions on Smart Grid*, vol. 5, no. 6, pp. 2777–2786, 2014.

[46] "MongoDB," www.mongodb.com, accessed: Jan 2018.

[47] IBM, "Why nosql? your database options in the new non-relational world," Tech. Rep., Mar 2015. [Online]. Available: https://cloudant.com/wp-content/uploads/Why_NoSQL_IBM_Cloudant.pdf

[48] M. Dewar, *Getting Started with D3: Creating Data-Driven Documents.* OReilly, 2012.

[49] D. Marr and E. Hildreth, "Theory of edge detection," *Proceedings of the Royal Society of London B: Biological Sciences*, vol. 207, no. 1167, pp. 187–217, 1980.

[50] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Commun. ACM*, vol. 60, no. 6, pp. 84–90, May 2017.

[51] R. S. Sutton and A. G. Barto, *Reinforcement learning: an introduction.* The MIT Press, 2012.

[52] L. Prechelt, "Early stopping - but when?" in *Neural Networks: Tricks of the Trade.* Springer-Verlag, 1997, pp. 55–69.

[53] P. Klein, J. Merckle, D. Benyoucef, and T. Bier, "Test bench and quality measures for non-intrusive load monitoring algorithms," *IECON 2013 - 39th Annual Conference of the IEEE Industrial Electronics Society*, 2013.