

LCDR Ethan Proper, Dr. Robert W. Cox, Dr. Steven B. Leeb, LCDR Keith Douglas, Jim Paris,
Warit Wichakool, Edward L. Foulks, LCDR Richard Jones, LT Perry Branch,
LT Ashley Fuller, LCDR Jeremy Leghorn, LCDR Greg Elkins

Field Demonstration of a Real-Time Non-Intrusive Monitoring System for Condition-Based Maintenance

ABSTRACT

The performance of important electrical loads on mission critical systems like warships or off-shore platforms is often tracked by dedicated monitoring equipment. Individual monitoring of each load is expensive and risky. Expense occurs because of the need for individual sensors and sensor wiring for every load of interest. Reliability is compromised because detected failures or fault conditions might legitimately be due to load failure, but might also be due to errors or failure in the sensor network or recording instruments. The power distribution network on a warship could be pressed into “dual-use” service, providing not only power distribution but also a diagnostic monitoring capability based on observations of the way in which loads draw power from the distribution service. This paper describes field tests of a prototype system that monitors multiple loads using existing electrical wiring. Initial results are presented from a device that monitors a small collection of motors and two other devices that monitor an entire engine room.

INTRODUCTION

In the modern world we are surrounded by sophisticated networking tools that make it easy to send and receive information. Examples include Ethernet, Bluetooth, 802.11b, and cellular networks. Unfortunately, the process of “feeding” and “clearing” a network, that is, of gathering and analyzing the data, remains expensive in many applications. For example, the information available to an engineering officer about a propulsion plant or any other engineering plant is generally directly proportional to the complexity and size of the installed sensor array.

During the past year, we have developed new hardware and software that makes it possible to construct a nearly sensor-less system for

monitoring the condition of mission critical loads. Our approach relies on electrical data (i.e. current and voltage) that is collected at a central point in a power distribution system. Our field-tested device is referred to as a non-intrusive load monitor (NILM), and it can identify electrical loads from aggregate current measurements and then perform diagnostics on the identified loads.

This paper provides a description of the automated load-identification and diagnostic system that is used to simultaneously monitor multiple loads in shipboard systems. The paper begins with a discussion of the potential benefits of non-intrusive load monitoring over more traditional approaches requiring large sensor networks. The paper then describes previous work conducted aboard ships. The third section describes a software package named *ginzu* that automates the load-identification and fault-detection processes. The paper then provides field examples from a NILM that monitors multiple loads in a waste-removal system. Finally, the paper concludes with results from an initial field study in which two NILMs monitored an entire engine room.

NON-INTRUSIVE LOAD MONITORING VERSUS TRADITIONAL APPROACHES

The development of high bandwidth networks has made an old dilemma increasingly more apparent: although networking makes it easy and inexpensive to obtain information from remote sensors, useful information can only be gathered by a potentially expensive and intrusive sensor array. Although mass production may ultimately reduce sensor cost, especially for solid-state or technologically advanced micro-electromechanical sensors, installation and

analysis will likely remain expensive. The overall reliability of a monitoring system with many sensors may be reduced in comparison to a system with relatively fewer sensors. The utility of data collected with a monitoring system is critically dependent on the ability to perform relevant and fast analysis of the collected data. More sensors may provide more potentially useful information, but at increased cost and increased burden in collating and correlating relevant observations.

On combat vessels, modern propulsion plant monitoring systems, for example, rely on hundreds of sensors arrayed throughout the main machinery space. Although these sensor networks enable increased levels of automation, they are costly to install and to maintain. As these networks grow to include more sensors, there is a corresponding drop in the overall reliability of the monitoring system.

Fortunately, the growing reliance on electrically actuated systems provides a new opportunity to reduce sensor count. The basis for this claim lies in the fact that electrical currents contain significant information about the physical condition of individual loads. A device that monitors aggregate current at a central location can then disaggregate and track the behavior of multiple downstream components.

The Non-intrusive Load Monitor (NILM) is a system that can determine the operating schedule of electrical loads in a target system using centralized measurements (Leeb 1995, Shaw 2008). In contrast to other systems, the NILM reduces sensor cost by using relatively few sensors. The NILM disaggregates and reports the operation of individual electrical loads like lights and motors using only measurements of the voltage and aggregate current at the service entry to an electrical panel.

Over the last decade we have been conducting an aggressive research program to develop the NILM as a nearly sensor-less platform for monitoring mission critical electromechanical loads on warships. Field experiments have been conducted on board two US Coast Guard Famous Class Cutters, the *USCGC Escanaba* and *USCGC Seneca*. We have also begun to examine monitoring possibilities for US Navy ships,

including the DDG-51 class destroyer, through experiments conducted at the Navy's Land-Based Engineering Site (LBES). Until recently, most of these experiments have focused on particular engineering subsystems. Also, the experiments have typically not involved real-time reporting of information to crew members. Results are presented in several publications including Cox (2006, 2007) and Mitchell (2007).

During the past year, we have developed and tested new hardware and software for nonintrusive load monitoring. We have tested the ability of our new monitoring system to provide useful information while underway, augmenting the observations traditionally made by a watchstander. In some cases we have provided new information for which no sensor had been previously installed. We currently use the NILM to monitor small collections of electrical loads, but we have initiated studies that consider how many electrical loads a NILM can successfully monitor on a shipboard power system. Our ultimate goal is to develop a practical lower bound on the power changes that can be effectively detected by a NILM installed at the switchboard level.

POWER SYSTEM MONITORING OVERVIEW

Power system monitoring is an exciting approach for creating an inexpensive, highly capable "black-box" for monitoring the performance of critical shipboard systems. With remarkably little installation effort or expense, we have fielded a minimally intrusive power monitor that can reliably monitor and track diagnostic conditions for multiple devices. This NILM can be used to determine the need for maintenance, to identify fault conditions, to find power quality problems, to help reconfigure a power system after damage, and to provide reliable verification of load operation. Generally, the power distribution system can, with the proper signal processing and data analysis, be made to serve "dual-use." That is, it can simultaneously be used for its intended function of power delivery and as an information network for monitoring critical loads.

As previously noted, the NILM makes measurements of voltage and current solely at a single point in the electric utility service. It characterizes individual loads by their unique signatures of power drawn from the mains. A transient detection algorithm can identify when each load turns on and off, even when several do so nearly simultaneously. This monitoring can be performed with relatively little hardware: a Pentium-class computer, an A/D converter, and a single set of current and voltage sensors (see Figure 1).




Figure 1: Block diagram of NILM installation

When installing a NILM to monitor multiple loads on a ship or other target system, it first undergoes a training phase. During training, the NILM observes individual electrical transient events that occur during the operation of particular loads. Sample data that might be observed by a NILM is shown in Figures 2 and 3.

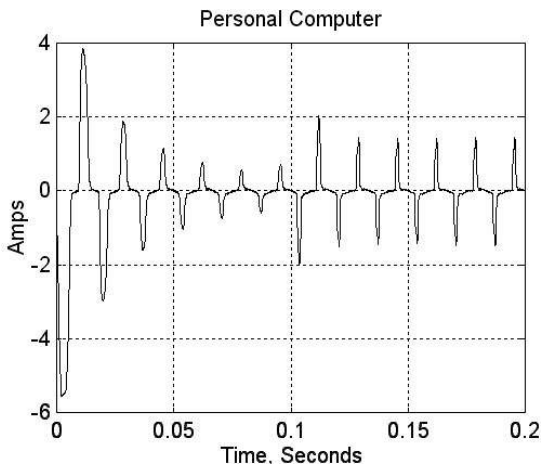


Figure 2: Turn-on transient of a power electronic load (personal computer)

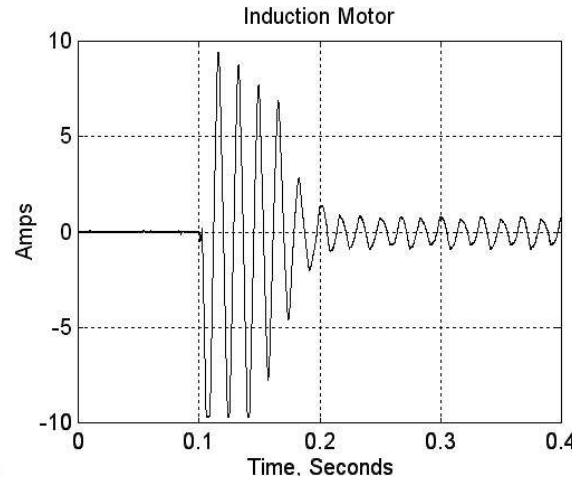


Figure 3: Turn-on transient of an induction motor pump

Figure 2 shows the turn-on transient of a personal computer. Figure 3 shows the turn-on transient of an induction motor turning a pump head. Each electrical load performs a different physical task, and each consumes power in a relatively unique way associated with its task. The personal computer shown in Figure 2, for example, is essentially a power electronic load that draws a distorted line current with substantial third harmonic content. This distortion occurs because the computer connects to the line through a full-wave rectifier. The current trace also shows the internal sequencing of load components, e.g., the step change associated with the activation of an internal electrical component approximately one tenth of second into the transient. The induction motor in Figure 3, on the other hand, draws a large pulse of current while accelerating the rotor, and then settles to a smaller steady state current demand. These transients serve as “fingerprints” that can be used to identify the operation of a particular type of load, even when several loads are operating at the same time.

In practice, the NILM examines or recognizes fingerprints by looking for known shapes in “spectral envelopes” or short-time estimates of the envelope of frequency content in the current waveform. An example is shown in Figure 4 for another induction motor. The top trace shows current versus time during the start transient. The solid line in the lower trace shows the component of current at the same frequency and phase as the line voltage, or a spectral envelope that

corresponds to real power in steady state. The “x” data points in the lower trace show a stored exemplar or fingerprint that the NILM uses to identify this particular load. .




Figure 4: Induction motor spectral envelope

The NILM examines the spectral envelopes of line frequency currents both in-phase and quadrature to the line voltage, as well as higher harmonics. The in-phase and quadrature fundamental frequency spectral envelopes are sometimes referred to as “real” and “reactive” envelopes. The induction motor, for example, would be characterized by traces of in-phase and quadrature line current. In the case of the power electronic load (personal computer), there would also be useful fingerprint information in spectral envelopes associated with third harmonic frequency (180 Hz on a 60 Hz utility).

We have engaged in research to explore the possibility of using a nonintrusive approach to diagnostic monitoring. This work has been directed at collecting, examining, and modeling data from field observations that we have collected onboard the *USCGC Escanaba* (Fig. 5) and the *USCGC Seneca*, as well as at the USN LBES facility. Seven key systems have been monitored onboard the Coast Guard Cutters, including auxiliary seawater (ASW) pumping, vacuum-assisted waste disposal systems (collection-hold-transfer or CHT), and reverse-osmosis water purification (RO). We have observed these systems both in-port and underway during operational cruising.



Figure 5: USCGC ESCANABA

The unique signatures presented by different classes of loads create an opportunity for diagnostic monitoring. Once it becomes possible to associate observed waveforms or segments with specific loads, it is possible to perform state and parameter estimation on the observed waves to track and trend diagnostic parameters for individual loads. For example, on the *USCGC Seneca* we have used the NILM to determine several important operating parameters of the ASW system. The auxiliary seawater system provides cooling for all heat loads onboard the cutter with the exception of those associated with main diesel engine cooling. Heat loads that are cooled by this system include the HVAC units, refrigerators, freezers, diesel engine air coolers, and diesel engine lube oil coolers.

ASW provides an excellent example of the diagnostic capabilities of the NILM. A coupling that connects the ASW pump motor to the pump head can fail, leaving the ship temporarily without cooling, a major mission complication. Field data showed that a high frequency “ripple” present in the spectral envelopes during pump transients increased as the coupling progressively failed. The NILM is able to track each start of the ASW pump and evaluate the condition of the coupling by computing a diagnostic metric from the observed spectral envelopes (Shaw 2008). The progress of a typical coupling failure is shown in Figure 6. The imminent failure of the coupling can be predicted three to five starts before ultimate failure using power line monitoring.




Figure 6: ASW coupling experiencing progressive failure.

REAL-TIME MONITORING

During the past year, we have worked to automate the recognition of load operation and diagnostic monitoring to make results available to the crew in real-time. To do so, we have developed a software package known as *ginzu* that eliminates the need for off-line analysis by a skilled observer. Initial tracking of load operation and diagnostic condition are now provided automatically by the NILM on-board ship. Furthermore, our field-tested systems have been installed at a central point that allows them to monitor multiple loads simultaneously.

The *ginzu* software application implements a detect-classify-verify loop that locates electrical load transients, identifies them using a decision-tree-based expert classifier, and then generates event files that contain relevant information. Additionally, the *ginzu* application provides streaming data to a graphical user interface known as the *Ginzu* Graphical User Interface (*GinzuUI*).

Classification Overview

In general, NILM classification methods have focused on identifying system-specific events based on signal characteristics. The methods implemented in the *ginzu* classification software compare the shape characteristics of a transient to shape characteristics of known events. Specifically, the shape characteristics are defined as (1) the relative steady state power change

across the transient event index and (2) the shape of the spectral envelope during the transient. The comparisons are aided by continuously tracking the state of the system (i.e. the running status of the known motors and other electric components in the system) and limiting the classification decisions to only those permitted by the associated finite state diagram of possible operating conditions.

Figure 7 shows a simplified flow diagram of *ginzu's* program logic. The algorithm initializes by loading a 10 second data window consisting of relevant spectral envelopes, e.g., corresponding to in-phase and quadrature (“real” or “P” and “reactive” or “Q”) components of current. This window is passed to a detection algorithm that locates vector indices where rapid changes in the envelopes have occurred. These indexes represent system transients and are candidates for classification.

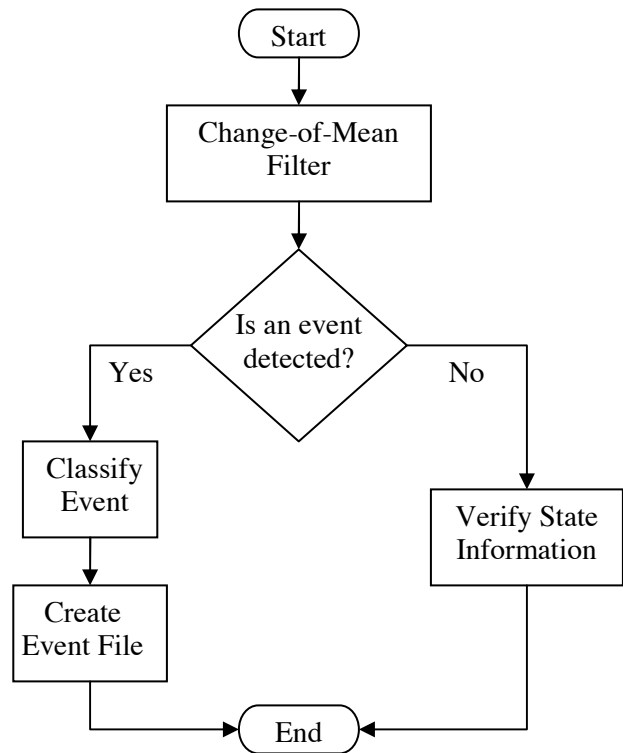


Figure 7: Program flow diagram. Shows detect-classify-verify logic used for incoming power.

Once an event has been detected, the classifier may be called. The classifier implements a hierarchy of classification decisions to make a ‘best guess’ based on the relative power levels around the event, the state of the system prior to the event, and if possible, the correlation between the shape of the power signal during the event and the shape of a known library event.

On the other hand, if a rapid power change is not detected, a *state verification and correction* function is called. This function attempts to verify that the current power levels are consistent with what the *ginzu* algorithm anticipates them to be based on current system state. The algorithm then reads one additional period of data from the input buffer; this data is inserted into the P/Q buffer and the old data is discarded. This new P/Q window is then passed to the detect-classify-verify loop, and the cycle is repeated.

The following sections provide an overview of the main components of the *ginzu* software along with a brief description of the *GinzUI* application. For in-depth discussion refer to Proper (2008).

Event Detection

The preprocessor located upstream in the program flow of the *ginzu* software provides spectral envelopes for fundamental and higher harmonic content at a sample rate of 120 Hz. Therefore, the ten second data windows form several 1200 index arrays. For example, one array contains “real” power and another “reactive” power. The 1200 index power array is passed to the detection algorithm where rapid power changes are located. This is accomplished by using a change-of-mean filter that calculates the difference between the original power signal and the output of a low pass filter. The result is a processed signal that only contains rapid power changes. Ultimately, these power changes are compared to pre-determined detection thresholds (based on the monitored system’s characteristics). The output of the comparator is an index of ‘Event Detections’.

Classification Techniques

The *ginzu* software algorithm recognizes events by examining changes in both steady-state consumption levels and also transient shape.

Steady-State Power Change

When individual loads are cycled within the system (i.e. pump on/off), they produce a corresponding change in the real envelope (ΔP) and reactive envelope (ΔQ) and possibly other spectral envelopes. These changes can be used as a simple classifier to identify loads.

It is advantageous to look for changes in steady-state levels in as many spectral envelopes as contain useful information. This is illustrated in Figure 8, which shows steady-state power levels after turn-on and turn-off of two loads, a computer and a lamp. The top plot in Figure 8 shows changes in Q versus P. The middle plot shows the change in third harmonic content versus P. The bottom plot shows change in Q versus third harmonic content. Notice that in the P-Q space (top graph), the two loads are essentially indistinguishable. They both turn on and consume approximately 150 watts in steady state, with essentially no reactive power. Similarly, they turn off with a -150 watt steady-state change as expected. Observations summarized in the top graph alone could not be used to differentiate the operation of the two loads.

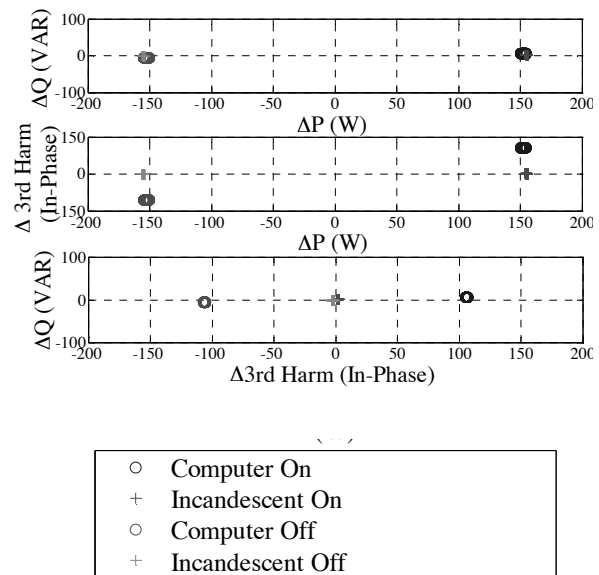


Figure 8: Steady-state changes for turn-on and turn-off of a computer and a lamp (see text).

The computer, however, draws a third harmonic current, distinguishing it from the lamp. The middle and bottom traces in Figure 8 show that this difference makes it easy to detect the operation of the computer with respect to the lamp in an information space that includes as many useful spectral envelopes as can be reliably recorded.

Transients

The *ginzu* algorithm also classifies individual loads based on distinctive load transient shapes. Overall transient profiles tend to be preserved even in loads that use active waveshaping or power factor correction. Most loads observed in the field have repeatable transient profiles, or at least sections of the transient profile that are repeatable.

Transient-based recognition permits near-real-time identification of load operation, especially turn-on events. Transients are identified by matching events in the incoming aggregate power stream to previously defined transient signatures, or “exemplars.” Exemplars can be determined, for example, by a one-time direct observation of the device in question, or by previous training in the laboratory. Pre-training has proven to be a reasonable approach for very repeatable loads that show up in large quantities, such as fluorescent lamp ballasts. The exemplar may be comprised of multiple parts for loads whose transients have a number of distinct sections. Each section of the exemplar can be shifted in time and offset to match incoming transient data. In addition, an overall gain may be applied to all sections of the exemplar to achieve a better fit. Each event detected is compared to the full set of exemplars by using a least squares criterion to select the appropriate shifts and gains. The match with the lowest residual norm per number of points is then compared to a threshold. If the fit is good enough, the event is classified as a match to the exemplar. If not, the event is left unclassified.

Correct classification of overlapping transients is possible using properly designed exemplars. Fingerprint traces provide positive identification of specific events occurring during system operation. By comparing the shape of the transients to known system events, a numerical

score can be assigned to grade the degree of similarity of the two signals; this score is known as the correlation score. It is derived using the method of least squares in the *ginzu* algorithm. This method is discussed in detail in Lee (2003) and Proper (2008).

State Verification

If no transients are detected within a given window, the classifier does not need to be called. *Ginzu* uses this opportunity to verify the current state of operation. This is accomplished by calculating the average power level for a ten second window and its standard deviation. These values are used in a state verification function to perform various checks and correct the state status if needed.

GinzUI

The *GinzUI* application provides the interface between the event file and the NILM user. Figure 9 illustrates the front-end display used on board the *USCGC Escanaba*. The primary functions of *GinzUI* are:

- To continuously check the user interface directory for newly created *event files*.
- To read *event file* contents and move the *event files* to an archive directory.
- To perform diagnostics on *event file* data and alert the user if a diagnostic has failed.
- To allow the user to graphically view *event file* contents.

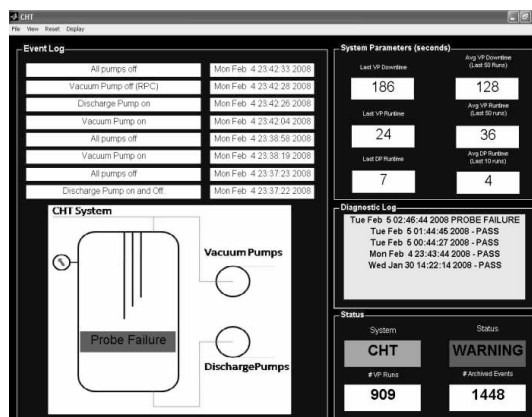


Figure 9: *USCGC Escanaba* CHT GUI

CHT TEST SYSTEM

To demonstrate the ability of *ginzu* to monitor multiple loads, we installed a NILM at the service entry to the power panel supplying the Shipboard Waste Collection and Disposal System (CHT) onboard the *USCGC Escanaba*. The CHT system represents a common shipboard auxiliary system used to transfer sewage from installed heads to a sanitary tank where it is pumped overboard.

NILM has been monitoring the CHT system since 2003 and various problems have been detected and classified through the application of NILM signal analysis. The CHT system operation and performance has been detailed in previous research (Mosman 2006, Piber 2007).

Figure 10 shows the prototype installation aboard the *Escanaba*. The NILM system with *ginzu* software runs completely on the touch tablet computer shown in Figure 10, installed next to the power panel serving the CHT system. The crew can interact with the NILM through the touch screen, both through diagnostic reports and also by requesting additional data and analysis from the NILM.



Figure 10: Real-time NILM monitoring CHT on the *USCGC Escanaba*.

A simplified schematic of the CHT system is provided in Figure 11. It consists of a 360 gallon sewage collection tank, which collects drains from eighteen vacuum toilets, two urinal lift valves, one urinal non-lift valve, and the ship's garbage

grinder. There are four pumps associated with the system including two vacuum pumps and two discharge pumps. A number of other ancillary single-phase loads (i.e. lamps and small motors) are also installed.

The vacuum pumps are each rated at 1.5 horsepower and connected upstream to the top of the collection tank and downstream to the vacuum seal tank. Their function is to maintain the necessary vacuum on the system for proper operation. If pressure falls below 14 in-Hg, one of the pumps automatically turns on to increase vacuum within the tank. The pumps alternate operation in order to minimize wear. If pressure falls below 12 in-Hg, both pumps are energized to restore proper vacuum pressure. Pumps are automatically secured when pressure reaches 18 in-Hg.

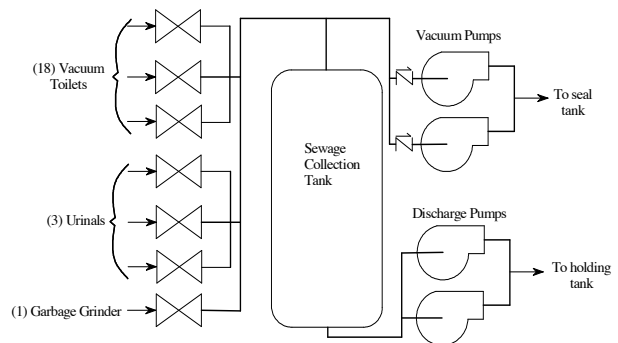


Figure 11: Simplified diagram showing key components of the CHT system.

CHT System State Recognition

As previously mentioned, the system states can be defined by measuring the real power usage. This approach was applied to the CHT system. The final allowable states that were defined for the CHT system are shown in Figure 12. By identifying the most likely state transitions and tracking these states, the classifier algorithm can be tuned so that the most likely transitions are given additional consideration.

One additional note on Figure 12 is that it includes the most common transitions from state to state. In other words, if an ON event is detected while both vacuum pumps are already running, the event cannot be a vacuum pump ON.

As stated in the previous section, the state information can be combined with the power change information to create accurate classifiers. Consider an event where the post-event power was zero. Any loads known to be operating before the change should now be classified as OFF.

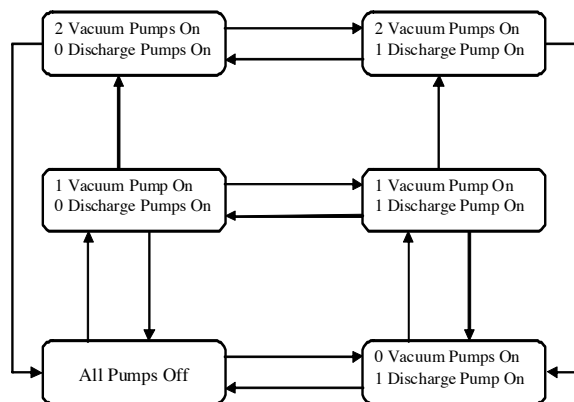


Figure 12: Finite state model for the CHT system.

CHT Diagnostic Package

A rudimentary diagnostic package was included in *GinzUI* to provide real-time detection of CHT system faults that had been observed by Mosman (2006) and Piber (2007). These failures include clogs in gauge lines and/or priming orifices, tank level probe failures, and system leaks. When a system fault is detected, a comment is printed to the diagnostic log and the log turns yellow to indicate an abnormal condition. Additionally, a line is printed to a text file that contains the description of the failure and the time of the detection.

FIELD RESULTS

The deployment of a complete diagnostic NILM system on the *USCGC Escanaba* took place in January of 2008. The installed system consisted of the *ginzu* event classification software and the *GinzUI* diagnostic GUI running on an IBM T60 tablet-style laptop. The *Escanaba* implementation was intended to provide the crew with immediate functionality with little or no special training. Ultimately, the package needed to be user friendly, stable, and capable of providing an automatic restart function to recover from any

scenario where the program execution was stopped.

Summary of Results

During its maiden cruise, the *ginzu* classifier automatically identified over 50,000 transients. Additionally, the associated GUI performed real-time diagnostics that identified the occurrence of:

- A continuous fault in the CHT tank probe level indication circuit causing over one thousand cycling discharge pump events.
- Two situations where the vacuum pumps were not reaching their normal power level due to probable clogging.

The CHT Classifier was assessed by randomly sampling and manually classifying 1500 CHT events from the *Escanaba* dataset. In these 1500 events, 62 events were determined to be misclassifications. This equates to 95.9% accuracy. The 62 events provide a representative sample of common classified errors.

Automatic Diagnosis of Probe Failure

Approximately five hours after the installation of the NILM, the diagnostics module alarmed that a probe failure had been detected based on excessive cycling of the discharge pumps. This particular fault had been previously diagnosed in research conducted by Piber (2007). It was discovered that the CHT electrical controller was not functioning properly following a modification from design specifications. Figure 13 shows a screen shot from the actual system. Notice that there is a short burst of discharge pump operation. This behavior is not typical and is characteristic of this particular fault.

Automatic Diagnosis of Vacuum Pump Clogging

Two vacuum pump clogging events were diagnosed during the three month patrol. When a clog occurs in the pump suction or priming line, its associated vacuum pump fails to add vacuum to the system, thus forcing the other pump to be energized. Figure 14 shows an actual screen shot recorded while this fault was in place. Notice both the expected 'double start' sequence and the lower power absorbed by the first pump.

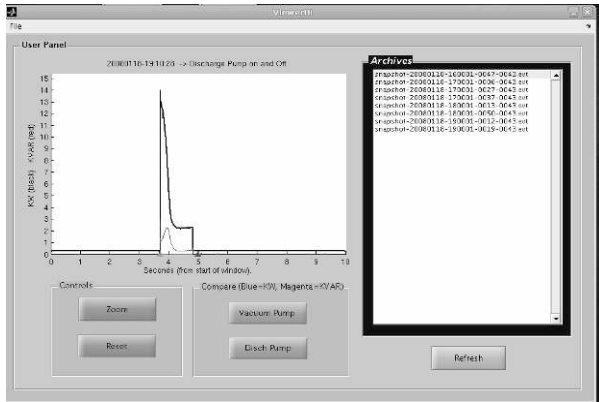


Figure 13: GinzUI showing eight cycling discharge pump runs.

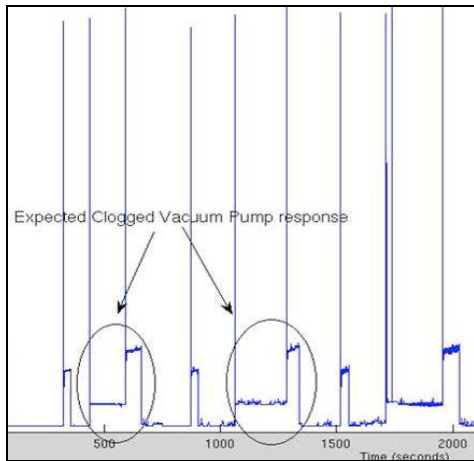


Figure 14: Vacuum pump clogging. The second pump starts to aid the first pump.

TESTING THE LIMITS OF NONINTRUSIVENESS

The success of the results aboard the USCG Cutters has encouraged us to continue our effort to develop and expand the role of the NILM as a platform for condition-based monitoring and maintenance. Ideally, as few NILMs as possible would be installed onboard a ship. However, there is a natural trade-off between intrusiveness and detection accuracy. Best certainty in event detection would be provided by individually monitoring every load of interest. Maximum nonintrusiveness, on the other hand, might involve one or a very small number of NILMs monitoring a large collection of loads on the power system, with a concomitant risk of mis-identification.

We are currently examining data from our field work at the USN LBES facility, where several NILM units have been installed. Some of these devices monitor individual loads, while two others are recording aggregate data at the two switchboards providing power to the entire engine room. Figure 15 and Table 1 present some initial results. The top trace in Figure 15 is an aggregate power signal recorded at a switchboard while the LBES was in use for crew certification. The bottom trace shows the power absorbed by the 2A lube oil service pump (LOSP) over the same time period. The letters correspond to positively identified events, and a summary is provided in Table 1. Notice that many major auxiliary loads are detectable and that the transient shape at the load and upstream are quite similar.

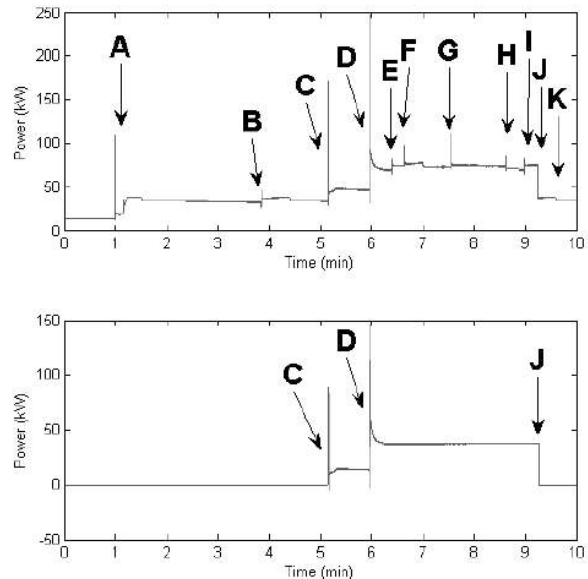


Figure 15: Aggregate power recorded during a crew certification at the DDG-51 LBES facility (top) and power absorbed by the 2A LOSP (bottom).

Results such as those shown in Figure 15 are very encouraging because they indicate that many critical loads can be detected with little additional design effort. We are currently comparing the dynamic range and resolution of the monitoring system at various points in the power system to develop a more rigorous design metric. This metric will indicate for the designer the likely tradeoff in identification accuracy versus the number of NILMs needed in a shipboard power system.

Table 1: Summary of events shown in Figure 15.

Event	Description
A	LPAC #2 Started
B	LPAC #2 cycling
C	2A LO SP Started in Slow Speed
D	2A LO SP Shifted to Fast Speed
E	2A FO SP Started in Slow Speed
F	LPAC #2 cycling
G	2A FO SP Shifted to Fast Speed
H	2A FO SP Shifted to Slow Speed
I	2A FO SP Secured & LPAC #2 cycling
J	2A LO SP Secured
K	LPAC #2 cycling

CONCLUSION

We have demonstrated the first real-time NILM that has provided diagnostic information in near real-time to a serving military crew. The commissioning process for the NILM requires system knowledge, but is not onerous. It has been shown that the NILM is capable of performing as a stand-alone diagnostic tool. The performance on the *USCGC Escanaba* indicates that the NILM is a maturing and capable technology that could work well supporting ICAS and other condition-based maintenance efforts in the USN and USCGC.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the assistance and support provided by the crews of the *USCGC Escanaba* and *USCGC Seneca*. Their patience has been indispensable. The author's gratefully acknowledge the support and technical assistance of NEMOmetrics Corporation, and Dr. John Rodriguez, Dr. Rod. Hinman, and Dr. Neil Lupton. This work was supported in part by the Grainger Foundation, the National Science Foundation, NASA, the United States Coast Guard, the U.S. Navy through ONR's ESRDC

program, NAVSEA, and the University of North Carolina.

REFERENCES

- Cox, R., J. Mosman, T. McKay, S. Leeb, and T. McCoy. 2006. Diagnostic indicators for shipboard cycling systems using non-intrusive load monitoring. In *Proc. ASNE Day 2006*, June, Arlington, VA.
- Cox, R., G. Mitchell, P. Bennet, M. Piber, J. Paris, W. Wichakool, and S. Leeb. 2007. Improving shipboard maintenance practices using non-intrusive load monitoring. In *Proc. ASNE Intelligent Ships Symposium VII*, May, Philadelphia, PA.
- Lee, K. 2003. Electric load information system based on non-intrusive power monitoring. Ph.D. diss., Massachusetts Institute of Technology, Cambridge.
- Leeb, S., S. Shaw, and J. Kirtley. 1995. Transient event detection in spectral envelope estimates for nonintrusive load monitoring. *IEEE Trans. on Power Delivery* 10: 1200–1210.
- Mitchell, G., R. Cox, J. Paris, and S. Leeb. 2007. Shipboard fluid system diagnostic indicators using non-intrusive load monitoring. In *Proc. ASNE Day 2007*, June, Arlington, VA.
- Proper, E. 2008. Automated Classification of Power Signals. S.M. thesis, Massachusetts Institute of Technology, Cambridge.
- Shaw, S., S. Leeb, L. Norford, and R. Cox. 2008. Nonintrusive load monitoring and diagnostics in power systems. *IEEE Trans. on Instrumentation and Measurement*. 57(7): 1445–1454.
-
- LCDR Ethan Proper** received the B.S. degree in Electrical Engineering from the University of Wisconsin in 1995 and the Naval Engineer and S.M. degrees from MIT in 2008. He is currently an Engineering Duty Officer in the United States Navy.

Dr. Robert Cox received the S.B., M.Eng., and Ph.D. degrees from the Massachusetts Institute of Technology (MIT) in 2001, 2002, and 2006, respectively. He is currently Assistant Professor of Electrical and Computer Engineering at UNC Charlotte. His research is focused on the design, analysis, and maintenance of electrical actuators, power-electronic drives, analog instruments, and sensors.

Dr. Steven B. Leeb received the S.B., S.M., E.E. and Ph.D. degrees from MIT. He currently holds appointments as Professor of Electrical Engineering and Professor of Mechanical Engineering at MIT. He is concerned with the design, analysis, and development of all kinds of machinery with electrical actuators, sensors, and power electronic drives.

LCDR Keith Douglas, LCDR Jeremy Leghorn, and LCDR Greg Elkins are currently graduate

students in the Naval Construction and Engineering Program (Course 2N) at the Massachusetts Institute of Technology

James Paris and Warit Wichakool are pursuing doctoral degrees in Electrical Engineering at MIT.

Edward Lane Foulks is pursuing a doctoral degree at UNC Charlotte. He was previously an ETI nuclear reactor operator aboard the USS Theodore Roosevelt (CVN-71).

LCDR Richard Jones, LT Perry Branch, and LT Ashley Fuller are currently serving as Engineering Duty Officers in the US Navy. They are graduates of the Course 2N graduate program at MIT.