CDR William Greene, USN; LCDR J. Steven Ramsey, USN; Steven B. Leeb; LT Thomas DeNucci, USCG; James Paris; LT Mike Obar, USCG; Robert Cox; Christopher Laughman; CDR Timothy J. McCoy, USN

Non-Intrusive Monitoring for Condition-Based Maintenance

ABSTRACT

The Navy's Integrated Condition Assessment System (ICAS) is continuously evolving to meet the needs of today's fleet as well as the requirements of future ship designs. This paper describes a technology known as the non-intrusive load monitor that could support the evolution of ICAS and other condition-based maintenance (CBM) systems. The non-intrusive load monitor (NILM) [8, 9, 10] can be used in many instances to determine the state of electromechanical systems strictly from easily acquired electrical measurements. This paper explores the potential applicability of the NILM to naval vessels using field data collected on the *USCGC Seneca*.

INTRODUCTION

Monitoring of machinery systems has become increasingly important in meeting the rapidly changing requirements of U.S. warships. As the pressure to reduce manning on ships increases, so does the need to reduce organizational level maintenance. Increased operating tempos are requiring maintenance providers to make repairs faster and ensure that equipment operates reliably for longer periods.

ICAS

The Integrated Condition Assessment System uses a CBM approach to reduce the burden of preventive maintenance practices on crews while improving coordination with shore-based maintenance providers. A shipboard ICAS installation is typically made up of several Windows NT workstations, printers, portable (manual) and installed (automatic) data acquisition devices, and a CD-ROM tower of Navy logistic products, all linked by a fiber-optic Local Area Network. ICAS data is transmitted from ships to shore-based support facilities on a periodic basis, or immediately in the event of a key equipment failure. ICAS data potentially improves the efficiency of technical support visits and evaluation. It also potentially improves the accuracy of future work packages and class maintenance databases. ICAS is the U.S. Navy's "Program of Record" for CBM and is currently installed on over 97 ships fleet wide [6].

In order to benefit from the use of CBM systems and smaller crews, it is imperative that ships have large sensor networks to provide information regarding component status. As a result, modern naval vessels are now equipped with vast and costly arrays of advanced sensors. Current ship classes such as the DDG-51 are adding significant numbers of sensors through retrofits, while some estimates indicate that the DDX will have as many as 250,000 sensors [2].

Non-Intrusive Monitoring

The NILM is a device that can determine the operating schedule of all of the electrical loads in a target system strictly from measurements made at the electric utility service entry [8, 18]. For example, the NILM can disaggregate and report the operation of individual electrical loads such as lights and motors from measurements of voltage and current made only at the electric meter where utility service is provided to a building [9, 11]. It can identify the operation of electromechanical devices in an automobile from measurements made only at the alternator [15]. The NILM is capable of performing this disaggregation even when many loads are operating simultaneously. We have begun to consider the application of the NILM to ship-board power systems [3, 7].

The NILM is in many ways an ideal entry point for measuring and collating useful information about any system that uses electromechanical devices. It requires a bare minimum of installed sensors, reducing expense and potentially enhancing system reliability. Because the NILM can associate observed electrical waveforms with the operation of particular loads, it is possible to exploit modern state and parameter estimation algorithms to verify the operation and "health" of electromechanical loads [11, 12, 13, 16]. The NILM can also monitor the operation of the electrical distribution system itself, identifying situations in which two or more otherwise healthy loads interfere with each other's operation through voltage waveform distortion or power quality problems [10, 14].

Challenges Facing Future Diagnostic Systems

In the next few years, ICAS may have the ability to predict a fault on a particular machine or electromechanical load. With that ability, the ICAS could alert the ship's control system of the impending failure so that the load in question could be secured and an alternate unit could be brought on line [4, 6]. One of the many challenges facing the Navy is the rapidly increasing number of sensors required to achieve this vision. The data communications wiring required for machinery monitoring systems "makes up a large part of the overall system complexity, cost and weight." One estimate for non-military industrial wiring of this kind is $5-10 per foot [5]. In addition to installation costs, cables are costly to maintain and increase the footprint of a sensor system: "They are vulnerable to damage and need to be removed and re-run whenever equipment needs moving, replacement or maintenance" [1]. Finally, the amount of power required for a network comprised of tens or hundreds of thousands of sensors is likely to be significant.

As the number of shipboard sensors grows, the issues of cable cost, size, weight, maintenance, and power demand are magnified. The Navy's current DDG-51 Class already has 1,342,000 feet of cables to support electrical power distribution, communications and sensors [19]. By introducing the NILM into an engineering space or a similar substantial section of a ship's electrical distribution system, it may be possible to mitigate the numerous economic and implementation issues associated with a large sensor network. Using data from field experiments conducted on-board the *Seneca*, the following sections explore the feasibility of using the NILM on operational vessels to reduce the overall number of sensors.

SHIPBOARD APPLICATIONS OF THE NILM

A complete NILM system consists primarily of commercial-off-the-shelf (COTS) hardware. A shipboard monitoring system consists of a customized NEMA-type enclosure to house the measuring transducers for voltage and current at a point on the (typically three phase) power distribution system. The remainder of the NILM includes a Pentium-class computer, keyboard, monitor, data acquisition card (either PCI or USB) and an uninterruptible power supply. The computer executes the custom NILM signal processing software that disaggregates individual load events from the aggregate current and voltage data. Figure 1 shows a NILM on-board the *Seneca*.

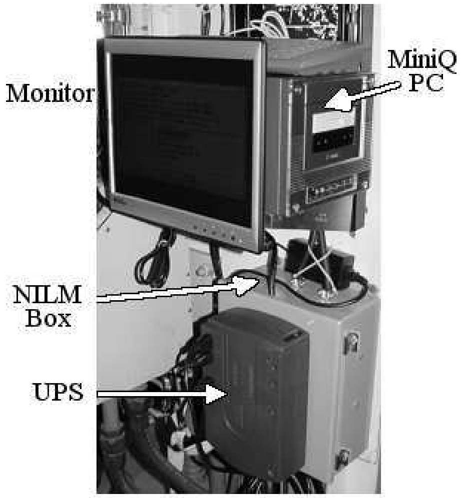

FIGURE 1: NILM installation on the *Seneca*

Several NILM systems have been installed onboard the *USCGC Seneca*, each monitoring a relatively small collection of loads. For example, the NILM system shown in Figure 1 monitors a collection of four motors (two vacuum pumps and two transfer pumps) for the waste-water handling system. The NILM systems installed on the *Seneca* intentionally monitor collections of loads that are small and relatively easy for the NILM to recognize and disaggregate. We made the decision to monitor small sets of loads during field testing in order to focus on the development of diagnostic indicators for particular loads of interest to us and the crew. Examples of results from *Seneca* are presented in Figures 2 through 5 below, including data from the auxiliary sea water (ASW) pumping system for heat loads, the waste-water vacuum pumps, and the rudder hydraulic steering gear.

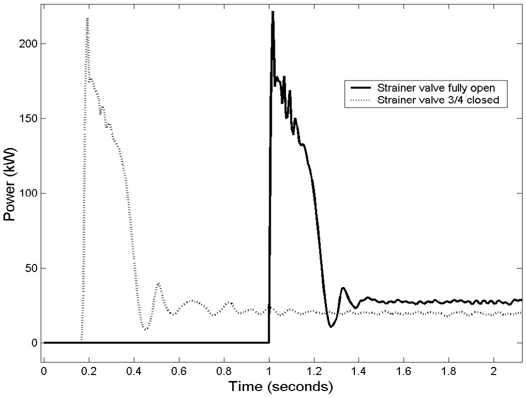

FIGURE 2: ASW pump starts with inlet flow restriction

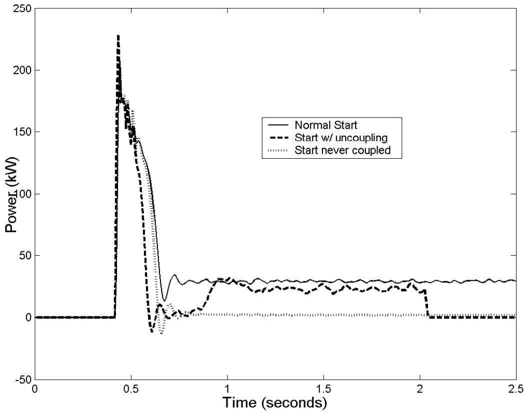

FIGURE 3: ASW pump starts for various levels of motor and pump coupling

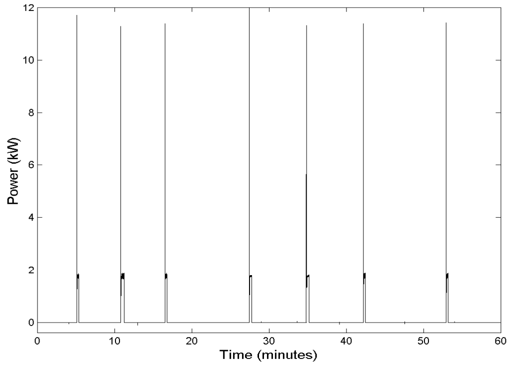

FIGURE 4: Sewage system vacuum pump transients

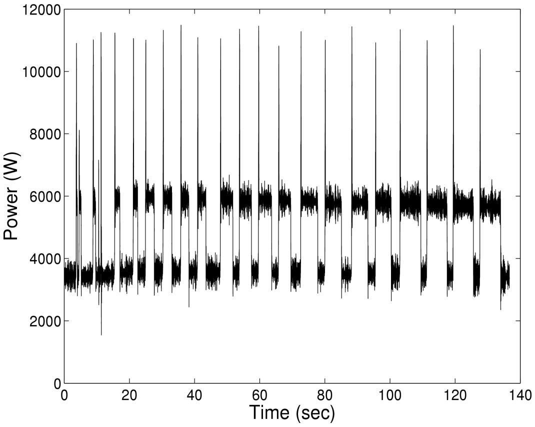

FIGURE 5: Steering pump transients

The NILM measures voltage and current on the power distribution system at the point where it is installed. From this data, the NILM continuously computes spectral envelopes that correspond to the harmonic content of the current waveform with respect to the phase of the voltage waveform. Harmonic current content at the line-voltage frequency corresponds, in steady-state operation, to conventional definitions of real and reactive power flow [8]. The graphs in Figures 2 through 5 illustrate the short-time harmonic content corresponding to real power flow during operation of the indicated loads. The NILM also computes reactive power flow and higher harmonic content, although these traces are not shown in the figures.

Notice in Figures 2 through 5 that different loads performing different physical tasks exhibit characteristic transient shapes that can be used as "fingerprints" to recognize the operation of particular loads. During installation at a site, the NILM is trained to recognize these fingerprints. Following training, the transient event detection algorithm can recognize and identify these fingerprints or templates even when many loads are turning on and off at the same time.

In some cases, the figures show both normal and pathological signatures, e.g., Figures 2 and 3. Figure 3, for example, shows start-up transients for an ASW pump during normal operation and also during an impending failure of the motor-to-pump head mechanical coupling. The NILM can use the differences between "healthy" and "pathological" transient observations to recognize impending failures or maintenance needs. References [3, 7-13, 16, 17] discuss the adaptation of the NILM to perform diagnostic monitoring using a variety of approaches to signal processing and state and parameter trending.

To ameliorate the issues associated with the rapidly increasing number of sensors on naval vessels, a single NILM must be able to perform load detection and CBM for more than just a few loads. If a single NILM can monitor the loads in a reasonably sized engineering space, then significant economies might be achievable. For example, a NILM capable of monitoring a significant collection of loads might be collocated with the Multi-function Monitor (MFM) electrical protection equipment installed on many ships. In this case, each NILM would find a natural location on the ship's power distribution system, and would monitor a significant collection of loads.

INVESTIGATING THE POTENTIAL OF THE NILM

In order to monitor more loads, a NILM must be installed further upstream in the electrical power distribution network. As this is done, the NILM is said to become less "intrusive." Despite the obvious benefits presented by moving the NILM further from the component level, there is a trade-off involved. Specifically, as the NILM moves further upstream, it becomes increasingly difficult to identify the operation of individual loads from the measured aggregate current. A natural question arising from this situation is the following: "How 'non-intrusive' can the NILM be while still providing useful information about the operation of individual loads?"

The extent to which the NILM can be truly non-intrusive raises a direct trade-off between monitoring hardware expense and signal processing effort. That is, if several components are monitored by one NILM, sensor numbers will be reduced but signal processing demands will likely increase. To investigate the plausibility of monitoring multiple components with a single NILM, a simulation of a machinery space was developed using field data from the *Seneca*.

Development of a Machinery Space Simulation

Several NILM systems were employed on the *Seneca* to collect focused data on individual ship systems, including the ASW pumps, vacuum pumps, and steering gear hydraulic pumps. This data was collected with the NILM positioned close to small groups of loads of interest in order to search carefully for diagnostic indicators that could be used for NILM CBM. In practice, a NILM might ideally be located in a ship's engineering control center to monitor a larger collection of loads. Ideally, the NILM would not simply disaggregate and identify the operation of individual loads, but also recognize key diagnostic indicators determined through close examination of the loads during the research phase.

For the "close" examination that has occurred during our early field tests, each NILM is configured with an input sensing range for current and voltage designed to take full advantage of the NILM's analog-to-digital (ADC) conversion resolution. Load data from these experiments, therefore, comes with a "scaling factor" of watts per ADC count tailored for those loads. Some of our observed component power ratings and scaling factors for the real power spectral envelope are shown in Table 1. The differences in scaling factors are a logical consequence of the very different sizes of the loads of interest, ranging from the smallest power consumer, a sewage vacuum pump, to the largest consumer, one of the ASW pumps. The vacuum pumps can be monitored with a higher input signal gain on the NILM data acquisition system, providing signals with the smallest watts per count resolution in our experiments. Conversely, the NILM monitoring the ASW pumps is configured with the lowest input signal gain, yielding the largest watts per ADC count in our field work.

In practice, a genuinely "non-intrusive" NILM will employ a current sensor that is a compromise designed to permit observation of the largest power transients of interest while still providing the best resolution possible for the smallest transients of interest. To study the challenges that would be faced by a single NILM monitoring an aggregate current feeding all of the loads in Table 1, raw data observations of every load were put on a common scale. That is, base data from individual load observations could be summed with high accuracy off-line in Matlab. This summed data can be rescaled and digitally quantized to emulate the effect of observing the actual aggregate stream with a particular ADC front-end sampling the combined current signal. In other words, given the fine observations of small collections of loads made by several NILMs, it is possible to use this data to assemble one stream with the information content that would have been produced by a single NILM monitoring all of the loads of interest.

TABLE 1: Power ratings and scaling factors

Component	Rating (Hp)	Scaling Factor (watts/count)
#1 ASW pump	40	7.11
#2 ASW pump	40	7.11
#2 Vent Fan	15	0.642
#1 Sewage Vacuum Pump	1.5	0.619
#1 Sewage Discharge Pump	2	0.619
#1 Steering Pump	15	6.31
#2 Steering Pump	15	6.31

In this section, summed and re-quantized data will be used to determine the plausibility of using the NILM to monitor complex combinations of loads from a single point. For these experiments, the aggregate data was re-quantized using a simple scheme assuming that the real-power spectral envelope would be represented with 12 bits. In fact, this is an overly draconian simplification. The NILM is able to employ significant signal processing and also 14-bit data conversion on the current and voltage measurements to enhance the input dynamic range. The simple 12-bit quantization in this section is a conservative choice.

As an example of this re-quantization procedure, consider a machinery room simulation using several of the components listed in Table 1. Evaluation of the component data reveals a maximum power that could be consumed by this collection of loads on an aggregate service, that is, a single utility feed for the collection of loads. A NILM monitoring this aggregate load would operate with a front-end ADC scaling that permitted full-range observation from zero watts to this maximum power level. To create a conservative, quantized data stream representing what would be seen on the aggregate service, waveforms representing the sum of power consumed by these loads would be divided by this maximum power, and scaled in a range between zero and 4095 (the maximum number that can be represented by a 12-bit ADC). Finally, the "floor" function in Matlab is applied to the data to account for quantization, thus eliminating decimal fractions and producing a hypothetical real power spectral envelope waveform consisting of integer values between 0 and 4095. A Matlab script helps perform the summing and rescaling, making it easy to conduct hypothetical studies using different collections of loads and assuming different operating schedules.

Waveform Recognition

With the ability to assemble realistic aggregate waveforms attributable to collections of loads of interest, the NILM can be easily tested off-line to determine the likelihood of successful load recognition. The transient event detector (TED) for the NILM has evolved, and is discussed in several different incarnations in [8], [11], and [17], for example. The full TED employed in the

NILM uses information from different spectral envelopes, and also at different points in time, in order to make a successful transient identification. This section reviews a vastly simplified approach to event detection that both illustrates some of the tools used in the TED and also provides quick, rough assessments of the likely success of the NILM in the anticipated aggregate load environment.

This simplified TED works as follows. Identifying a fast-varying section of an observed transient creates a "fingerprint" signature for a load of interest. Figure 6, for example, shows the start-up transient of one of the vacuum pumps onboard *Seneca*. A fingerprint template vector, *t*, consisting of N samples might be constructed by sampling a varying region of the appropriately scaled and quantized transient, derived between 3.45 and 3.6 minutes, for example, in Figure 6.

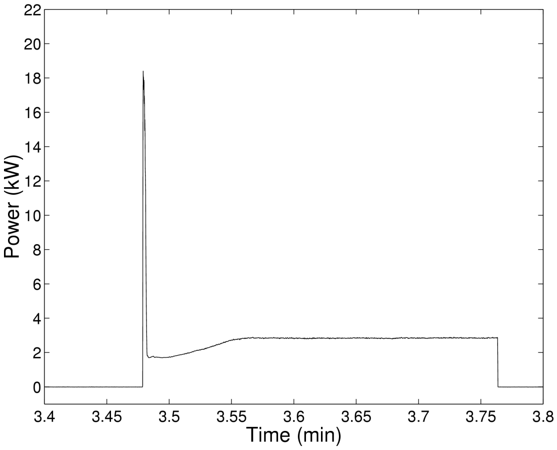

FIGURE 6: Vacuum pump start transient

The vector, *t*, consists of elements:

$$t[i], \quad i = 0 \ldots N-1.$$

An “ac-coupled” and amplitude normalized version of this vector, designated *tac*, can be computed as:

$$tac = \frac{\left[t - \frac{1}{N}\sum_{i=0}^{N-1} t[i]\right]}{\| \left[t - \frac{1}{N}\sum_{i=0}^{N-1} t[i]\right] \|^2}$$

A transversal filter can now be used with this template, *tac,* to search an incoming data stream for the vacuum pump “fingerprint” [8]. Such a transversal filter would have an impulse response corresponding to the time-reversed samples of the vector *tac*. When the filter response is convolved with an incoming data stream, the output of the filter will be unity whenever the original template points, *t*, appear in the data stream.

Essentially, this process computes the inner product between the ac-coupled and amplitude normalized template, *tac*, and a sliding window of data sampled from the incoming aggregate data stream. A unity output could indicate a perfect match. It could also indicate a window of data in the input stream with a unfortunately large norm and a different shape from the template. For this reason, a more sophisticated approach to detection is employed in the NILM TED. The simple approach described here illustrates a key computational component of the more sophisticated NILM TED, and also permits a quick check of the reasonability of deploying the NILM to monitor a large collection of loads.

Five templates were developed for component start-up transients from the quantized and scaled data (Table 2).

TABLE 2: Target component templates

Template	Component
t1	#1 Sewage Vacuum Pump
t2	#1 Sewage Discharge Pump
t3	#2 Vent Fan
t4	#1 Steering Pump
t5	#1 ASW pump

A comparison was made of these templates with respect to the original transient data for that load and also with respect to other load data windows used to develop templates. This comparison is presented in Table 3. When a template is compared to the entire signature from which it originated, the transversal filter output peaks at a match value of 1.0, as seen in the center diagonal of Table 3. When a template is compared with a signature of a different component, the result is a peak value other than 1.0. In reality, small variations in transients and noise from other loads make the case of a perfect match unlikely. Thus, it is wise to determine a range of peak values that can be considered to be a matching range. The values off the center diagonal in Table 3 provide some means for determining at least a crude set of boundaries. For example, these comparisons suggest that any peak value in the range from 0.9 to 1.1 could be considered a recognizable event.

A total of seven tests were conducted using combinations of components from Table 1. The tests are grouped into three sets and are summarized in Table 4. The overall goal of these experiments was to determine if a target event could be detected against the background noise of the simulated engine room. The results of each test are described in the following section.

TABLE 3: Waveform match values

Transient Signature	Templates				
	t1	t2	t3	t4	t5
Vacuum	1.0	0.652	0.352	1.543	0.533
Discharge	0.864	1.0	0.201	1.751	0.052
Vent Fan	1.267	0.6	1.0	2.073	0.065
Steering	0.415	0.339	0.221	1.0	0.021
ASW	16.5	13.26	5.48	28.41	1.0

Tests

The first set of five tests was performed in order to evaluate the ability of the basic TED to identify a single target transient against a steady-state background signal. To implement these tests, one representative transient for each of the five loads listed in Table 2 was added to a synthesized steady-state background signal at 3 different points in time. Figure 7 shows one such composite signal. A similar signal was created for each of the five loads listed in Table 2, and each of these was passed to the basic TED. The background signal used for each test was the same, and it was comprised of steady-state signatures produced by each of the following loads: the #1 ASW pump, the #2 ASW pump, the #2 Vent Fan with clean filters, and the #2 Vent Fan with blocked filters. Table 5 lists the match values for each test case, and Figure 8 shows the TED output waveform for a series of 3 sewage vacuum pump starts.

TABLE 4: Test table

Test Set	Test Number	Description
1	1	Recognition of vacuum pump starts against steady-state background
	2	Recognition of discharge pump starts against steady-state background
	3	Recognition of vent fan starts against steady-state background
	4	Recognition of steering pump starts against steady-state background
	5	Recognition of ASW pump starts against steady-state background
2	6	Recognition of steering pump start against multiple steering pump transients
3	7	Recognition of ASW pump start against multiple steady-state and transient events

The near-unity peak match values listed in Table 5 indicate that each of the loads can certainly be recognized individually while other loads are in steady operation. That is, a simple disaggregation of a load from a background stream of other loads operating in steady-state would be no problem in this hypothetical scenario.

A second set of numeric experiments was conducted in order to evaluate the ability of a particular template to identify the operation of its associated load from several different transient

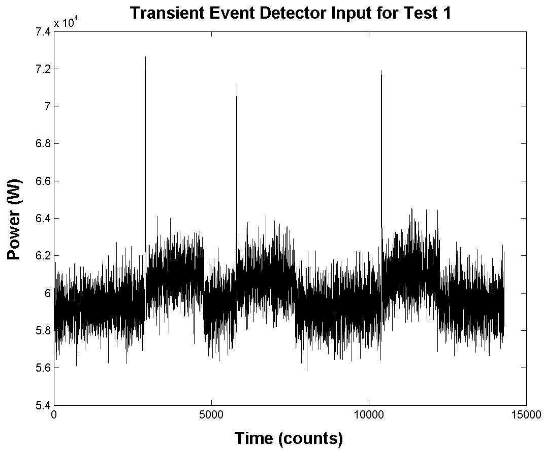

FIGURE 7: TED input for test 1 – Composite of 4 steady-state component signatures and 3 sewage vacuum pump transient signatures

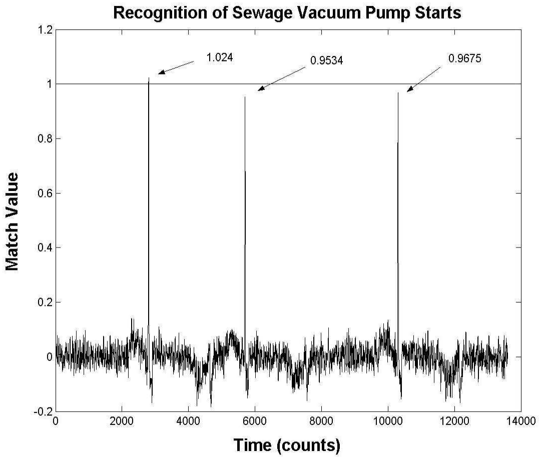

FIGURE 8: TED output for test 1 - Recognition of 3 sewage vacuum pump starts

events. In this case, instead of superimposing identical transients, the data used for this test were 16 different steering pump start transients recorded while the *Seneca's* rudder was fishtailing (i.e. the behavior shown in Figure 5). Although this test could have been performed with any of the loads listed in Table 2, the steering pump serves as an excellent example since our "real world" field data shows slight variability in the steering pump transient signature from start to start. Ultimately, this test will be repeated exhaustively for each load.

TABLE 5: Match values for tests 1-5 - Motor starts against steady-state background

Test/ Motor Start	Match Values			Mean
	Event #1	Event #2	Event #3	
1 Vacuum Pump	1.024	0.9534	0.9675	0.9816
2 Discharge Pump	1.011	1.035	0.9963	1.0141
3 Vent Fan	1.001	1.033	0.9767	1.0036
4 Steering Pump	0.9555	0.9761	1.0200	0.9839
5 ASW Pump	0.9991	1.001	0.9973	0.9991

Each of the 16 different transients used in the second set of tests was added to the same synthesized steady-state background signature used in the previous set of tests. Again, the component signals were processed using the engine room emulation and waveform recognition processes described previously. In this test, 12 of the 16 transients yielded peaks in the 0.9 to 1.1 match value range. The peaks of 4 of the 16 steering pump starting transients were slightly below the match bounds of 0.9 to 1.1. These four values are shown in Table 6.

TABLE 6: Match values exceeding bounds of 0.9-1.1 for test 6 – Recognition of steering pump starts during rudder fishtailing

Pump Run	Peak Value	Deviation from Match Band
2	0.8306	0.0694
6	0.8820	0.0180
7	0.8880	0.0120
11	0.8488	0.0512

In this case, the simplistic approach to transient recognition is slightly stressed by the variability in the pump transients and the background signals. The match deviations are small, however, and the more sophisticated TED used in the full NILM would be able to recognize even the four transients referred to in Table 6. In fact, a template constructed not from a single observation but from the average of several training observations of the steering pump, would allow even the simple detection scheme to work in the anticipated match range with success.

A final set of tests was performed in order to evaluate the ability of one template, the sewage vacuum pump template t1, to identify its associated transient against a background of both steady-state and transient component signatures. In this case, the sewage vacuum, sewage discharge, vent fan and ASW start-up transient signatures used to develop Table 3 were each separately superimposed on the synthesized steady-state signature used previously. The simple transient event detector easily determined that the only peak value in the matching range was due to the vacuum pump transient. Table 7 lists the peak match value for each transient.

TABLE 7: TED output for test 7 - Recognition of sewage vacuum pump start against multiple steady-state and transient events

Load	Peak Value	Deviation from Match Band
Discharge Pump	2.430	1.330
ASW Pump	16.53	15.43
Vent Fan Motor	1.280	0.180
Vacuum Pump	0.9653	zero

Test Summary

The results of the first set of tests showed that start-up transients, even from a small component such as the 1.5 Hp sewage vacuum pump, consistently generate a match value inside the selected detection range of 0.9 to 1.1. The findings of the second set of experiments indicate that a single template will recognize many (12 of 16) similar transients generated by the same component. These results further indicate that the steering pump template would recognize all of the similar transients if the threshold were to be expanded slightly (by 0.0694) to include the outliers listed in Table 6. An examination of the column labeled t4 in Table 3 suggests that such a small change would not be likely to result in any false detections. The results of the final test demonstrate the use of a single template (vacuum pump) to recognize a component start-up against a more complex background of both steady-state and transient signatures. In this case the smallest component (the sewage vacuum pump) was still recognized.

CONCLUSION

The data analysis in this paper provides the first indications that the NILM with a full transient event detector could successfully monitor large collections of loads on a warship. A typical trade-off is also apparent. The most minimal monitoring installation might attempt to use one or two NILMs at the generating points on a ship to track all loads. This provides the cheapest arrangement in terms of installation effort and hardware expense. It also places the highest possible burden on the TED software, and plainly could lead to misidentifications. Alternatively, individual monitoring of every load provides accurate information (to the extent that the large network of sensors are all working) at the expense of substantial sensor installation and maintenance costs. In the coming year we hope to continue field tests and refine the understanding of just how “non-intrusive” the NILM can be while still providing acceptable performance.

Present machinery monitoring programs are laying the foundation for future systems that will one day provide a fully integrated ship control system. NILM is a sensor technology that has the potential to provide a great deal of operational and diagnostic information when in a “stand-alone” mode or when tied into existing machinery-monitoring systems such as ICAS. The testing conducted for this paper indicates that NILM also has the potential to help reduce the number of sensors in current and future warships by monitoring multiple components simultaneously. This in turn could greatly reduce the installation and maintenance costs associated with machinery monitoring. Once the NILM is able to associate observed waveform features with the operation of particular loads, it may be possible to compute diagnostic indicators and use the NILM as a platform for condition-based monitoring.

REFERENCES

[1] Brooks, T. and Lee, K. "IEEE 1451 Smart Wireless Machinery Monitoring and Control for Naval Vessels," *13th Ship Control Symposium, Orlando FL*, 2003.

[2] Davidson Instruments webpage. Avalable: http://www.davidson-instruments.com/news. July 10, 2002.

[3] J.Steven Ramsey, et.al., "Shipboard Applications of Non-Intrusive Load Monitoring," *ASNE Symposium, Jacksonville, FL*, February 2005.

[4] DiUlio, M., Savage, C. and Schneider, E. "Taking the Integrated Condition Assessment System into the Year 2010," *13th Ship Control Symposium, Orlando FL*, 2003.

[5] Harmon, R. and Clark, D. "Wireless Machinery Monitoring Systems for Shipboard and Power Plant Applications," *13th Ship Control Symposium, Orlando FL*, 2003.

[6] Naval Surface Warfare Center, Carderock Division, Philadelphia. Integrated Condition Assessment System webpage. Avalable: <https://icas.navsses.navy.mil>. August 2004.

[7] J. Steven Ramsey, "Shipboard Applications of Non-Intrusive Load Monitoring", MIT NSEE/S.M. EECS thesis, June 2004.

[8] S.B. Leeb, S. R. Shaw, and J. L. Kirtley, Jr. 1995. "Transient Event Detection in Spectral Envelope Estimates for Noninstrusive Load Monitoring," *IEEE Transactions on Power Delivery* Vol. 10 No. 3, pp. 1200-1210.

[9] Norford, L. K. and S. B. Leeb. 1996. "Nonintrusive Electrical Load Monitoring." *Energy and Buildings*, Vol. 24, pp. 51-64.

[10] Abler, C., R Lepard, S. Shaw, D. Luo, S. Leeb, and L. Norford. 1998. "Instrumentation for High-Performance Nonintrusive Electrical Load Monitoring." ASME *J. Solar Energy Engineering*.

[11] K.D. Lee, S.B. Leeb, L.K. Norford, P. Armstrong, J. Holloway and S.R. Shaw, Estimation of Variable Speed Drive Power Consumption from Harmonic Content, accepted for publication, *IEEE Transactions on Energy Conversion.*

[12] S.R. Shaw, M. Keppler, S.B. Leeb, Pre-Estimation for Better Initial Guesses, *IEEE Transactions on Instrumentation and Measurement*, Volume 53, No. 3, June 2004, pp. 762-769.

[13] S.R. Shaw, S.B. Leeb, Identification of Induction Motor Parameters from Transient Stator Current Measurements, *IEEE Transactions on Industrial Electronics*, Volume 46, No. 1, February 1999, pp. 139-149.

[14] S.R. Shaw, R.F. Lepard, S.B. Leeb, and C.R. Laughman, A Power Quality Prediction System, *IEEE Transactions on Industrial Electronics*, Volume 47, No. 3, June 2000, pp. 511-517.

[15] Shaw, S.R., D. Evangelista, S.B. Leeb, and C.R. Laughman, Nonintrusive Load Monitoring and Identification in an Automotive Environment, *Proceedings of ELECTRIMACS 1999*, Lisbon, Portugal, pp. 199-204, September 1999.

[16] P. Armstrong, "Model Identification with Application to Building Control and Fault Detection", Building Technology, Massachusetts Institute of Technology, Cambridge, MA, 1993.

[17] Shaw, S.R., "System Identification Techniques and Modeling for Nonintrusive Load Diagnostics," Ph.D. Dissertation, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, February, 2000.

[18] C.R. Laughman, K. Lee, R. Cox, S. Shaw, S.B. Leeb, and L. Norford, "Power Signature Analysis," IEEE Power and Energy Magazine, March/April 2003, pp. 56-63.

[19] M. Wickenheiser, "”Huge Ballet' Goes into Details of Building a Ship," Maine Sunday Telegraph, 11 April 2004, pg 7A.

ACKNOWLEDGEMENTS

This research was supported by a grant from the Office of Naval Research under the ONR Control Challenge Project. Data from the *USCGC Seneca* was made available from the ONR Electric Ship Research and Development Consortium project at MIT. The Grainger Foundation provided essential additional support. The authors gratefully acknowledge the valuable advice and support of Mr. Daniel Devine, John Amy, Katherine Drew, Sharon Beerman-Curtin, Terry Ericsen, and Professors Steven R. Shaw, James L. Kirtley, and George Verghese.

Authors

Commander Bill Greene is an Engineering Duty Officer in the U.S. Navy specializing in ship maintenance. He received his bachelor’s in Nuclear Engineering and an MBA from the University of Florida. Cdr. Greene has served on the three submarines; USS Guitarro (SSN-665), USS West Virginia (SSBN-736), and the USS Santa Fe (SSN-763). He recently completed a tour of duty at Puget Sound Naval Shipyard where he served as a Project Deputy on a submarine refueling availability. Currently he is pursuing graduate degrees in Naval Architecture/Marine Engineering and Mechanical Engineering at the Massachusetts Institute of Technology. Cdr. Bill Greene can be reached at wcgreene@mit.edu.

LCDR J. Steven Ramsey is an Engineering Duty Officer in the U.S. Navy. He earned a dual Master’s degree in Naval Architecture & Marine Engineering and Electrical Engineering from the Massachusetts Institute of Technology in 2004. He is currently serving at Norfolk Naval Shipyard in Portsmouth, Virginia.

Steven B. Leeb received S.B., S.M., E.E. and Ph.D. degrees from MIT. He has been a member of the MIT faculty in the Department of Electrical Engineering and Computer Science since 1993. He serves as an associate professor in the Laboratory for Electromagnetic and Electronic Systems. He is concerned with the design, analysis, development and maintenance processes for all kinds of machinery with electrical actuators, sensors or power-electronic drives.

LT Thomas DeNucci graduated from the U.S. Coast Guard Academy in 1998 with a B.S. in Electrical Engineering. He has served aboard *USCGC HEALY (WAGB-20)* and as a Port Engineer at Naval Engineering Support Unit (NESU) Seattle. He is currently a graduate student at the Massachusetts Institute of Technology pursuing a dual degree in Naval Architecture & Marine Engineering and Mechanical Engineering.

James Paris, Robert Cox, and Christopher Laughman are graduate students pursuing doctoral degrees at the Massachusetts Institute of Technology.

LT Mike Obar is a 1996 graduate of the United States Coast Guard Academy. Following assignment as a student engineer on *USCGC* Campbell, Lt. Obar served as a Port Engineer in Cleveland, OH. He graduated from the University of Michigan in 2001 with a degree in Naval Architecture and Marine Engineering. Currently, he is the engineer officer onboard *USCGC Seneca*. LT Obar is professional engineer licensed in the state of Michigan.

CDR Timothy J. McCoy, an Engineering Duty Officer, is presently assigned to MIT where he is an associate professor on the Ocean Engineering Faculty. He teaches the Naval Construction and Engineering curriculum to US Navy, US Coast Guard, International Navy and other students and conducts warship-related research. He holds a Ph.D., Naval Engineer and a master’s in Electrical Engineering from MIT and a bachelor’s degree in Mechanical Engineering from the University of Illinois. He has worked on the designs for the AOE-6, LPD-17 and DD(X) ship classes as well as the Integrated Power System (IPS) and Standard Monitoring & Control System (SMCS) development programs. He is a licensed Professional Engineer, a member of ASME, SNAME, USNI and a Senior Member of the IEEE.