VAMPIRE: Accessing a Life-Blood of Information for Maintenance and Damage Assessment

John Donnal, Uzoma Orji, Christopher Schantz, Jin Moon, Steven B. Leeb,
Jim Paris, Andrew Goshorn, Kevin Thomas, Jayme Dubinsky, Robert Cox

***Abstract*—Mission-critical machinery and associated test and maintenance equipment have not always been designed to speed repairs and minimize down-time. Life-cycle cost analysis and maintenance planning are not new ideas. However, we stand in a dawn of new sensing and communication technologies. Are we applying these technologies with best efficacy and return-on-investment for new hulls and retrofits? Are sensors revealing or a source of new maintenance problems? What are the best ways to process, present, and store data to provide the fastest decision assistance? Do our procedures and tools empower or hinder crews in speeding completion of maintenance and minimizing cost? This paper explores these questions through an example: vibration monitoring for USCG and USN vessels.**

I. Introduction

Profiles for shipboard maintenance costs arise from three essentially different engineering management strategies:

- Critical maintenance can occur on a scheduled or routine basis to attempt to ensure system availability, at the possible cost of an early "repair".
- Maintenance can occur when a system breaks, essentially deferring costs to a "day of reckoning" when the system is essentially guaranteed to be unavailable and therefore not mission capable.
- Watchstanders (including civilian or shore-side contractors) can attempt to optimize maintenance costs and system availability by working to predict needed repairs before systems become casualties.

The third option is appealing in the sense that it, at least intellectually, sounds thrifty. We work and think, and, if efficient and effective, achieve a desired degree of mission capability with well-reasoned expenses. However, all three approaches have both obvious and also subtle trade-offs between cost and mission availability. There is not necessarily a single correct approach, as the attractiveness of each philosophy will vary with the hull, the mission, and the expense in all resources (time, labor, training, budget, convenience, availability, etc.) inherent in the operation of different systems and ships. Even a specific analysis for a particular ship will vary with political, geographical, and mission-specific parameters, and may therefore change, possibly suddenly.

Complicating matters, we stand in the dawn of a technological change, where we expect electronics and communication systems to provide us with unprecedented levels of warning and situational awareness. These technological changes potentially introduce transformative effects for new hulls (such as USCG NSC or DDG-1000), and older hulls (such as USCG MEC or DDG-51) to some degree. Gathering useful information, however, may require the installation of an expensive and intrusive array of sensors. Can this burden be reduced?

Today's weapon system designer faces subtle questions about the nature and efficacy of installed sensors deployed for ensuring readiness and minimizing repair costs and operating expenses. Does the installation of a sensor add expense and risk through additional wiring? Do sensors expose metrics that could be part of a valuable maintenance decision assistant (MDA) or do they hide critical figures-of-merit associated with true system status? Our experience in field experiments on-board the USCGC *ESCANABA* during underway conditions indicate that, regardless of which of the three maintenance strategies might be deployed for maintaining mission critical systems, our choices of sensors and data processing strategy can directly improve, or equally directly limit, the "up time" and repair times for mission critical systems.

Time spent early in a system design implementing useful sensing and data reduction strategies for maintenance decision support may pay vast benefits for system availability and life-cycle cost.

We illustrate some of the possibilities for data gathering and extraction using vibration monitoring as an example. Both the USN and USCG employ vibration monitoring for determining system health for a variety

of rotating electric machine systems. The United States Coast Guard currently employs outside contractors to handle their vibration monitoring needs on a periodic basis [1], although USCG crews have conducted vibration monitoring in the past. There are interesting tradeoffs in "permanent" versus "occasional" monitoring. A preliminary study [2] showed that stator frame vibration changes as a function of the transducer position around the periphery of the frame. There is no firm guarantee that the transducer will be placed at the same location for each maintenance check, potentially undermining the longitudinal integrity of the readings. Outside experts incur expense. Contractor vibration checks are only conducted two to four times a year which may not catch problems.

Equally important in both USCG and USN approaches, where information may be gathered by contractors or by an *in situ* or permanently installed vibration sensor, respectively, is the measuring instrument itself and the nature of the test. Handheld vibration sensors condense information from a broad frequency range into a single "average" number or reading. This mixes signal energy and information from all vibration bands and eliminates valuable indicators. Underway, e.g., as part of USN ICAS [3], it may not be possible to cleanly distinguish vibration from a component versus coupled vibration from a neighboring machine, confounding diagnosis. In-port, as the USCG conducts vibration monitoring, it is possible to isolate the operation of a single or small number of machines. However, test conditions may not reflect realistic use conditions.

This paper demonstrates a work-in-progress for *in situ* vibration monitoring called vibration assessment monitoring point with integrated recovery of energy (VAMPIRE). The VAMPIRE prototype and preliminary signal processing experiments indicate features we may demand in future sensors: nonintrusive installation requiring no custom power or data wiring; maintenance-free, self-powered sensors; and scalable access to information, ranging flexibly from a "reduced" figure-of-merit to detailed time-series information permitting cross-correlation study with power consumption and operating condition. The approach and applications illustrated here, both in the laboratory and underway on board a USCG cutter, are intended to suggest possibilities beyond vibration monitoring. Our intent is to illustrate considerations for designing or retrofitting a system to minimize down-time, maximize the availability of "action-able" information, and reduce the expense and background maintenance associated with a sensor array. The approach can and will be extended to other environmental variables including temperature, humidity, power consumption, and corrosion, and for other systems besides rotating machines.

The proposed VAMPIRE can not only supply more detailed and rich information that can help diagnose types of imbalances, but can also potentially inform a real-time acoustic signature assessment as part of a tactical decision aide (TDA).

II. Signal Processing

Most data is susceptible to the possible requirement for interpretation without full context. That is, the watchstander or shore-side maintainer will likely never have all of the information necessary to interpret sensor data with complete confidence. Part of the pressure that moves the design of sensing systems for maintenance in one direction or another, depending on the maintenance philosophy, stems from opinions surrounding the process of wading through a river of data to reach a conclusion. Conclusions drawn from data collected with "hidden" correlations to unmeasured or ignored processes are potentially little better than numerology. Reduced numbers derived from sensors, e.g., averaged over time or frequency or both, essentially become "go/no-go" indicators.

We seek to provide a flexible approach for processing sensor data that scales required communication bandwidth, storage requirements, and computation with the needs and desire of the engineering duty team. A "boiled" number or figure-of-merit, like an average vibration reading, may be all that is desired for a quick look under time pressure. It will be shown in Section III that steady-state vibration can be viewed as a narrowband signal. The raw acceleration signal can be filtered by a passband filter and the steady-state vibration levels can be tracked over time as shown in Fig. 6 if an average vibration reading is all that is required.

More extensive information, e.g., energy in frequency bands, may be useful for understanding more complex pathologies. It will be shown in Fig. 8 that using this extensive information is essential in diagnosing underlying problems that cause abnormal levels of vibration. The ability to employ short-time frequency estimates to trade resolution and detail against computational complexity and data storage and transmission requirements has proven of great value in power system monitoring [4]. In the appendix to this paper, we review a related approach, the Hilbert transform, for vibration data, where, unlike the AC utility, there is no "system frequency" to synchronize the frequency analysis.

The Hilbert transform provides an easy implementation for developing amplitude envelope followers for

narrowband signals. The efficacy of using the Hilbert transform for such applications, e.g., vibration monitoring, is shown in the next section. We begin with a laboratory vibration analysis of two fans used for ventilation on USCG cutters. These fans illustrate the potential value of a more detailed look at vibration data correlated with operating condition (turn-on, turn-off, steady running, etc.).

III. Fan Vibration

Axial-flow fans can develop radial vibration due to imbalance. To better understand the conditions that lead to radial vibration, two ventilation fans, in the laboratory as shown in Fig. 1.

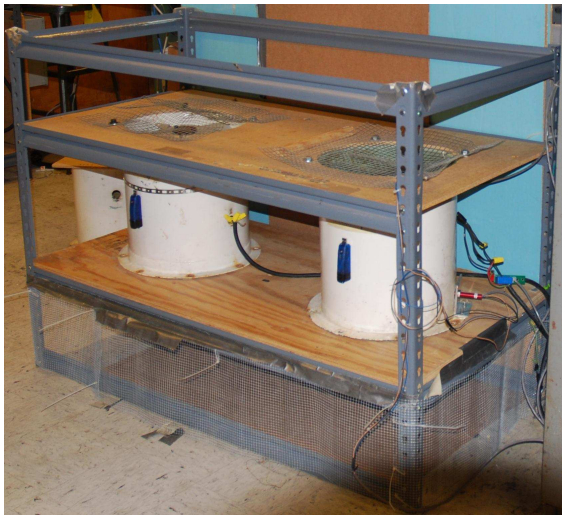

Fig. 1. Laboratory setup with two ventilation fans mounted to a steel frame.

The two motors are mechanically coupled through the steel frame. The motor on the left is the intake while the motor on the right is the exhaust. The following subsections detail the experiments conducted to test the effect of certain faults on the radial vibration.

A. Rotor Imbalance

A cause of radial vibration on the outer frame of the motor is rotor imbalance. This situation occurs in the field when fan blades are damaged or fouled or when bearings wear. An imbalance can be modeled as a weight attached to the rotor. With each rotation, this extra weight causes the body of the motor to displace which leads to vibration. In the lab, weights of different sizes were added to the hub of the fan and the radial acceleration was measured for each weight. Figure 2(a) shows the top of the fan in which no weight is added and Fig. 2(b) shows the fan that has been imbalanced.

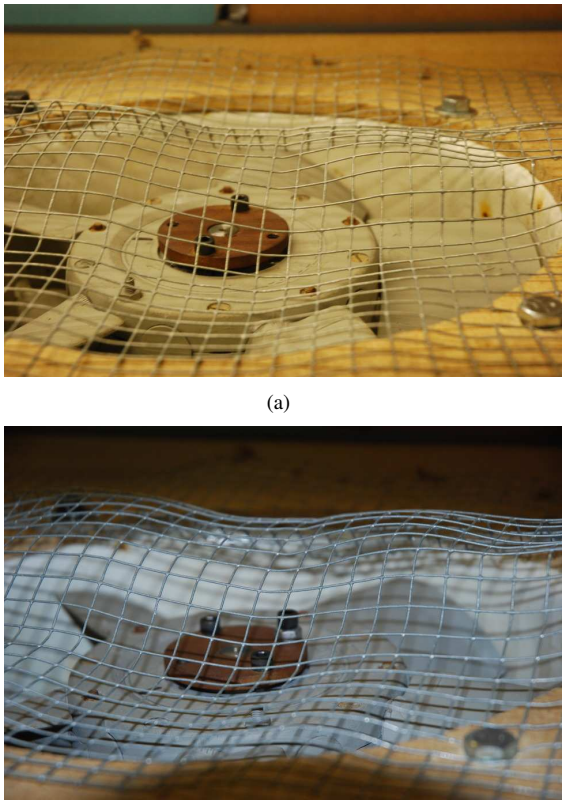

(a)

(b)

Fig. 2. In (a), a top view of the intake ventilation fan with no imbalance. In (b), a screw attached to the fan hub imbalances the motor causing radial vibration.

On the top of the hub, a screw, and a screw with a nut were added on different runs. Figure 3 shows the spectral content for the acceleration of the intake ventilation fan during five seconds of steady-state operation between 50 Hz and 60 Hz under baseline and faulty conditions. Almost all of the spectral content is located at the speed of the fan in rotations per second (rps). In this frequency band, the underlying signal can be approximated as

$$x(t) \approx A(t)\cos(\omega_c t),$$

where ω_c is the rotational speed of the fan and $A(t)$ is the time-varying amplitude of the sinusoid. During steady-state operation, $A(t)$ will remain approximately constant and will change only in the presence of a fault.

With high-resolution data and an appropriate windowing transform, e.g., the Hilbert transform in this case, we are not restricted to examining only steady-state information. The Hilbert transform can be applied to

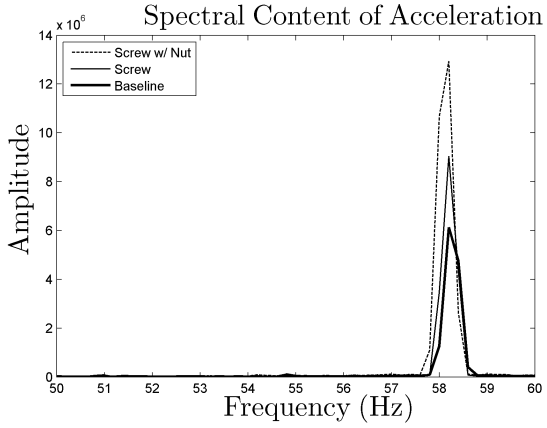

Fig. 3. Spectral content of the ventilation fan for five seconds of steady-state operation under baseline and faulty conditions between 50 and 60 Hz. All of the energy is centered around the frequency corresponding to the rotational speed of the fan. Approximating the energy in this frequency as a single sinusoid allows for the use of the Hilbert transform to track the envelope of the underlying sinusoid over time easily.

extract time-evolving content in small frequency bands of acceleration. For example, the acceleration levels in the 50 - 60 Hz frequency band for each run are shown in Fig. 4. The units of "acceleration" in this plot are "counts" from the analog-digital converter reading from the accelerometer, and are of course limited to a specific frequency band. So they do not correspond directly to the "g" reading on a hand-held accelerometer, but they are correlated with the vibration meter readings. With each increase in weight, the fan is more unbalanced and the amount of radial vibration increases as expected and can be tracked over time.

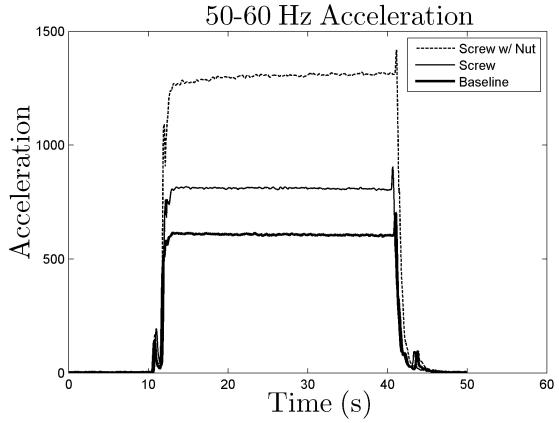

Fig. 4. Acceleration levels in the 50 - 60 Hz frequency bands for three different runs. The baseline test in which the fan has no imbalance gives the lowest amount of vibration. A second test was conducted with a screw added to the fan hub. The last test added a screw and a nut. The vibration levels increase accordingly.

It should be noted again that the passband energy used in this paper produce arbitrary units that cannot be converted to "g" or any other standard units of acceleration. However, as shown in Table I, the passband energy levels monotonically increase as the amount of eccentric weight on the fan hub is increased. The values from the handheld meter follows the same pattern.

TABLE I
NUMERICAL ENERGY LEVELS IN THE SECOND COLUMN FOR EACH OF THE TESTS CONDUCTED WITH THE INTAKE VENTILATION FAN. THESE READINGS ARE CONSISTENT WITH THOSE FROM A HANDHELD VIBRATION METER SHOWN IN THE LAST COLUMN.

	Passband Energy Level	Handheld Vibration Meter (g RMS)
Baseline	610	0.10
Screw	810	0.55
Screw & nut	1300	0.70

B. Loose Mounting

A second cause of radial vibration is from loose mounting. If a mounting or frame loosens over time, the mount can no longer dampen any vibration. This results in an increase in the radial vibration of the body of the machine. An improperly mounted machine can add noise which can be a source of distraction for the crew and a source of acoustic signature in the water. Mechanically, loose mounting also causes fatigue in the structure of the motor which can reduce the lifetime of the machine.

To emulate a loose mount in the laboratory, the mounting screws attached to the frame of the motor were loosened. Figure 5(a) shows a properly tightened screw which mounts the intake motor to the steel cage. The screw is loosened as shown in Fig. 5(b).

As the mounting screw was loosened, the amount of radial vibration increased as expected as shown in Fig. 6. These values are consistent from the values obtained from a handheld vibration meter. A summary of the values is shown in Table II.

TABLE II
THIS TABLE SHOWS THE NUMERICAL ENERGY LEVELS IN THE SECOND COLUMN FOR EACH OF THE TESTS CONDUCTED WITH THE INTAKE VENTILATION FAN. THESE READINGS ARE CONSISTENT WITH THOSE FROM A HANDHELD VIBRATION METER SHOWN IN THE LAST COLUMN.

	Passband Energy Level	Handheld Vibration Meter (g RMS)
Baseline	610	0.10
Loose	805	0.23

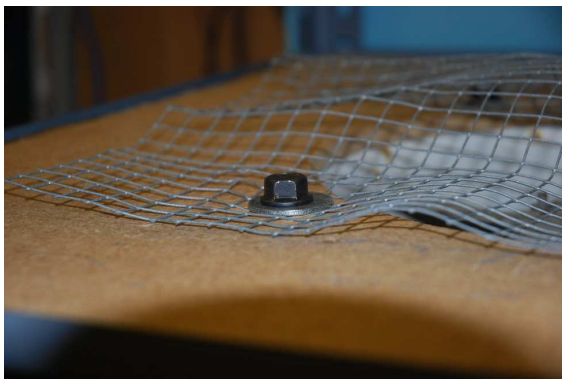

(a)

(b)

Fig. 5. An intake ventilation fan which is properly mounted to the steel frame is shown in (a). For testing purposes, the mounting screws are loosened as shown in (b).

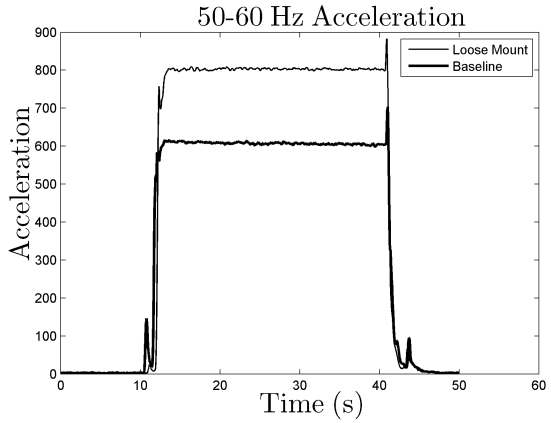

Fig. 6. The acceleration content in the 50-60 Hz frequency band for a baseline test is shown. A second run was conducted in which the mounting screws are loosened. The acceleration content is also shown. As expected, the readings show an increase in radial vibration.

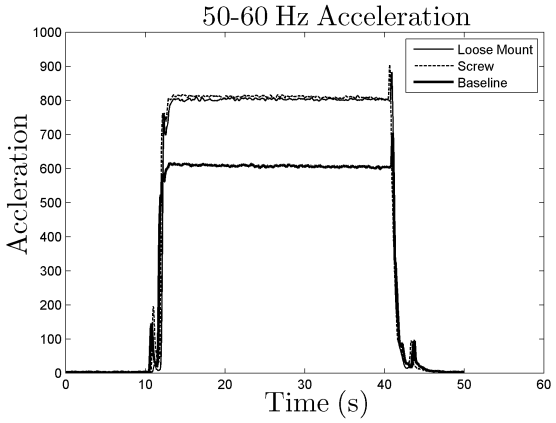

Fig. 7. Acceleration levels for the intake fan under normal ("baseline") conditions, with a rotor imbalance ("screw"), and with a loose mount. The faulty runs both show an increase in vibration. However, this information alone is insufficient to differentiate the two flaws. More analysis is needed to diagnose the underlying fault.

C. Differentiating Source of Vibration

Simple vibration readings can only indicate if there is a significant change in the amount of vibration. They cannot, however, diagnose the potential cause of the problem. In Fig. 7, the baseline for the intake motor is shown as the bottom traces. The top two traces depict the two faulty runs of the ventilation fans. In one run, the motor has a rotor imbalance by adding an eccentric weight to the fan hub. In the other run, the motor is loosely mounted to the steel frame. The vibration readings for the faulty runs are higher than that from the baseline but are indistinguishable from one another in this example. Diagnosing the problem between a rotor imbalance or a loose mount cannot be inferred from a vibration reading alone.

During turn-off transients of the ventilation fan, the spectral content of the fan's acceleration passes through the lower frequency bands as the fan slows down. When the fan is loosely mounted, the vibrational energy is dampened such that the envelope is indistinguishable from the baseline as show in Fig. 8. There is, however, no dampening effect with a rotor imbalance. The spectral content is higher than those from the baseline and the loose mounting runs. While vibration readings in the 50 - 60 Hz bandwidth give a simple indication as to the deviation from the baseline reading, vibration levels in the other frequency bands allow for a richer assessment. These methods for vibrational analysis were tested onboard the USCG *ESCANABA*.

IV. *ESCANABA* Data

Measurements discussed in this section were taken on board USCG *ESCANABA* (WMEC-907) shown in Fig. 9. Measurements were taken from two paired vacuum

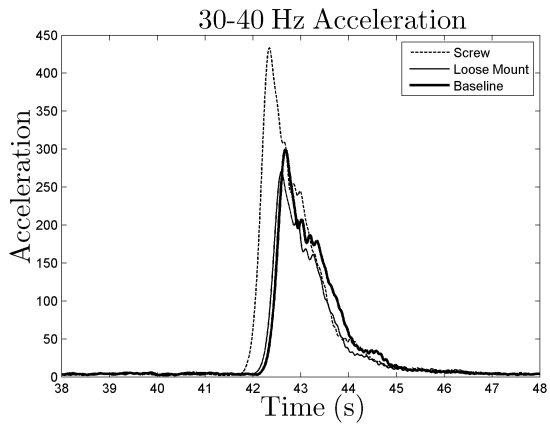

Fig. 8. Analysis of the 30-40 Hz frequency band is needed to differentiate between a rotor imbalance of a loose mount. During turn-off transients, there is more energy in the sub-bands when there is a rotor imbalance.

Fig. 9. USCG *ESCANABA* (WMEC-907). Measurements were taken from two paired vacuum pumps located above the ship's engine room.

pumps used to maintain pressure in the waste water tank while the ship was at port. The motors are located in the top compartment of the engine room. An image of these motors is shown in Fig. 10. The motor in the foreground will be referred to as Motor A. The motor in the background will be referred to as Motor B.

Observations were made during which an accelerometer was physically attached to each motor and the corresponding motor is turned on while the other motor remains off. The sample rate of the accelerometer is 160 Hz. All of the analysis is run against bands less than 60 Hz so the sample rate is higher than the Nyquist rate of 120 Hz.

Figure 11(a) shows the raw data of the acceleration for Motor A in a trial run. The motor is turned on at around 43 seconds and switched off at 141 seconds. The envelope in the 50-60 Hz frequency band of motor A is shown in Fig. 11(b). The raw data was filtered with

Fig. 10. Two vacuum pumps in the engine room on the *ESCANABA*. The motor in the foreground will be referred to as Motor A. The motor in the background will be referred to as Motor B.

a bandpass filter center around these frequencies. The amplitude of the Hilbert transform of this output signal is the envelope. This band was chosen because the nominal speed of the pump is 3420 rotation per minutes (rpm) or 57 rotations per second.

Ideally, the Hilbert transform is applied to narrow-band signals in which the amplitude changes slowly compared to the carrier frequency. In this example, this condition does not hold and there are oscillations. A moving average is computed on the signal to smooth out these oscillations During steady-state operation of the motor, the Hilbert transform can compute a metric that correlates directly with vibration. This value provides a baseline to which comparisons from future tests will be made. As with a handheld meter, this test condenses the information into a single number and serves as a reference point.

A second test was performed in which the the accelerometer is placed on motor A while motor A is off but motor B is running. The raw acceleration data is shown in Fig. 12(a).

Once again, the Hilbert transform is used to compute a numerical metric. It should be noted via visual inspection that this metric will be larger than that recorded when motor A was running. In other words, the accelerometer on motor A detects a higher vibration when another motor is on.

A purely automated analysis of the vibration readings (ignorant of operating schedule) from this experiment or from a vibration meter would inaccurately ascribe vibration readings to the wrong motor. In this instance, the

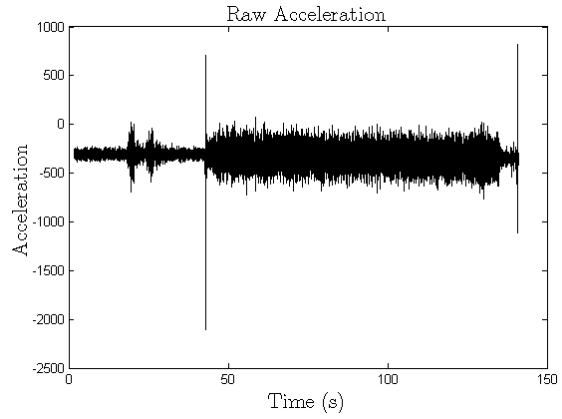

(a)

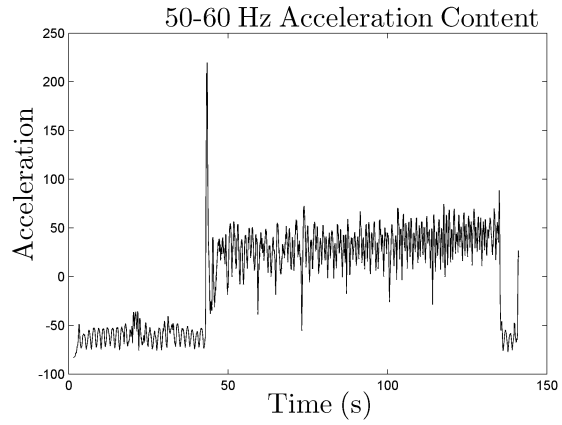

(b)

Fig. 11. Raw acceleration (a) of motor A when turned on. The acceleration content (b) takes the raw data and computes the envelope of the Hilbert transform of the data between 50 and 60 Hz. The acceleration content metric can be computed by taking a region when the motor is in steady-state operation.

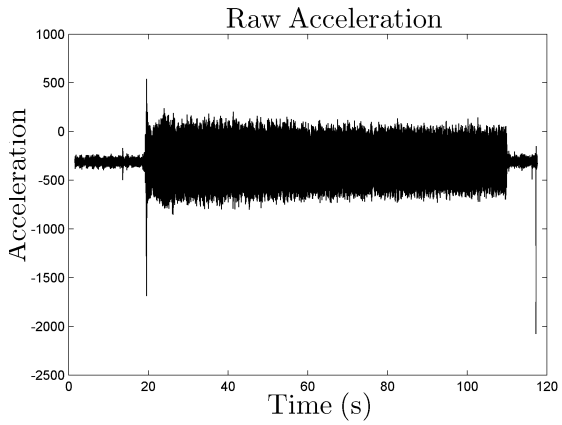

(a)

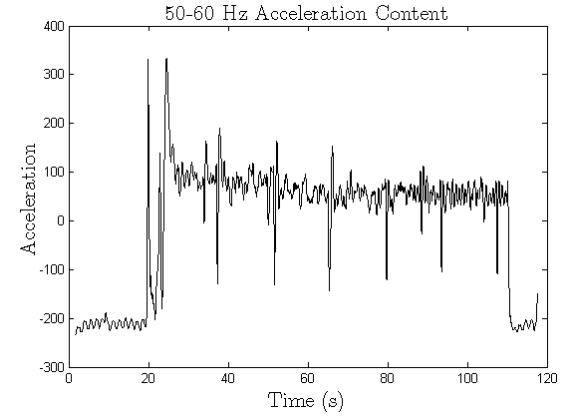

(b)

Fig. 12. In (a), raw acceleration of motor A when motor B is turned on. The amplitude of the signals is greater than that in Fig. 11(a). In (b), the acceleration content is similar to that in Fig. 11(b). These metrics can be misleading if no further analysis is done.

machines are mechanically coupled. A small vibration from Motor B can be enhanced through the coupling leading to a large vibration reading from an accelerometer placed on Motor A. These cross-interference effects cannot be inferred from boiled numbers.

However, analysis from the time-series data can show if and how two or more motors are mechanically coupled. Figure 13 shows the raw data of the acceleration on Motor A when both motors A and B are on. The interaction between the two motors is apparent as there is a beat frequency in the data.

If prior knowledge is known as to whether or not two motors are mechanically coupled, this information could be included in the analysis of vibration measurements. Furthermore, the fidelity of the data can be maintained depending on which motors are turned on. For instance,

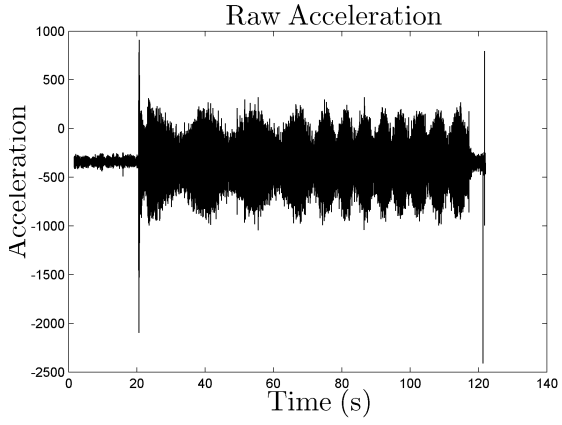

Fig. 13. The raw acceleration of motor A when both motors A and B are turned on. Visual inspection of the time series shows that there is coupling between the motors as there is a beat frequency present. This observation can be necessary in improving vibration diagnostics.

the data for the vibration of Motor A can be rejected if it is known that Motor B is also turned on. A non-intrusive load monitoring (NILM) environment provides a platform in which loads can be monitored by their turn-on and turn-off transients to keep track of which loads are on and off. A brief discussion of the NILM environment is described in the next section.

Not only can mechanically coupled motors provide false readings, but the actual sea state can also influence the data. Data of these motors were also taken when the *ESCANABA* ship was underway. Low-frequency acceleration caused by the rocking of the ship also add to the vibration readings. This phenomenon is readily seen in the raw acceleration in Fig. 14(a).

Before the motor turns on at the 13 second mark, a low-frequency oscillation caused by the ship's rocking is visible. A low pass filter removes the sea state from the rest of the data and its envelope is shown in of Fig. 14(b). All other frequencies add background noise which makes the handheld readings less reliable as shown in the second graph. A band pass filter around the rotational speed of the motor isolates the actual vibration reading from the motor shown in the Fig. 14(c). Processed time series data permits background information like sea state to be removed, revealing desired information.

V. VAMPIRE

Particularly for new designs, but for retrofit as well, there is a great attraction to a "dual-use" design philosophy. Dual-use seeks to piggy-back or squeeze extra functions from existing systems and structure. For example, the VAMPIRE acceleration monitor could be designed into the equipment to be monitored. We are working to design and develop a nonintrusive vibration monitor that requires no electrical connections. This new vibration monitor is a small toroidal transformer that self-powers from the magnetic fields around a wire feeding an electromechanical load of interest, e.g., a motor. This monitor could store data for later retrieval, or communicate wirelessly, or potentially communicate with a nonintrusive load monitor (NILM) collating data via a power line carrier modem. The system is maintenance free and easily installed as part of a new or retrofit motor enclosure cover. It could also measure temperature and other important diagnostic indicators.

Non-intrusive electrical monitoring has been described in [4, 5] among other publications. The NILM detects the operation of individual loads in an aggregate power measurement by preprocessing measured current and voltage waveforms to compute spectral envelopes [6]. Spectral envelopes are short-time averages of the line-locked harmonic content of a signal. These spectral

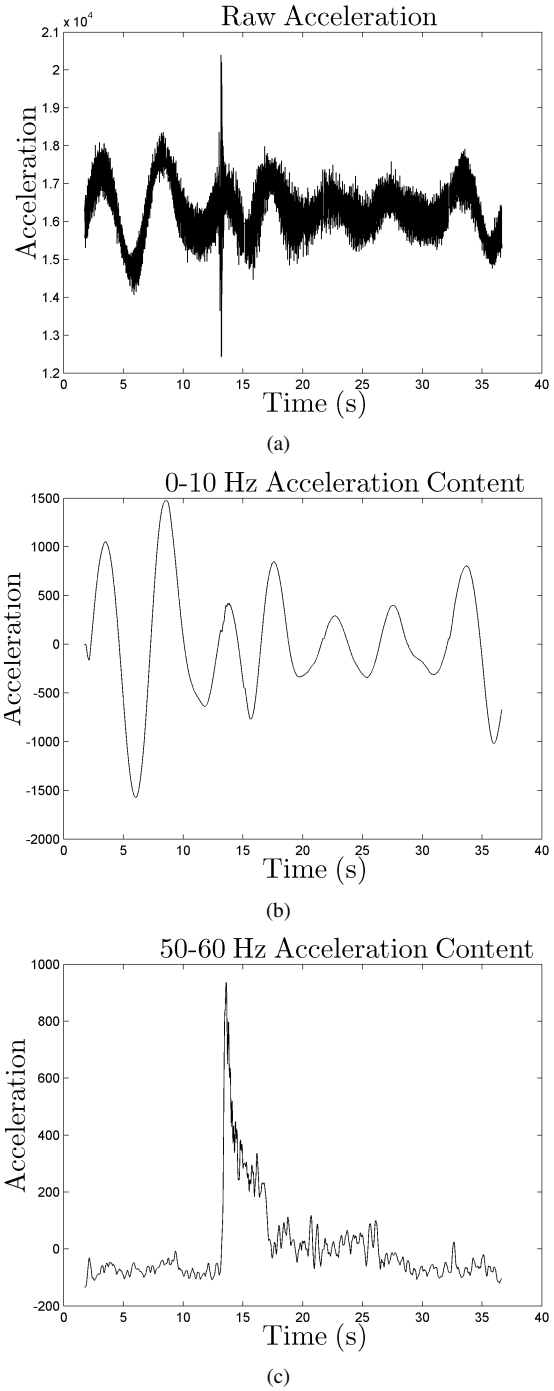

Fig. 14. Raw acceleration of motor A (a) while the ship was at sea with visible acceleration from the rocking of the ship. The decoupled low-frequency rocking is seen in (b). The actual vibration from motor A is shown in (c).

envelopes may be recognized as the coefficients of a time-varying Fourier series of the current waveform. For transient event detection on the ac utility, the time reference is locked to the line so that fundamental frequency spectral envelopes correspond to real and reactive power in steady-state. Higher spectral envelopes correspond to line frequency harmonic content. A high performance transient event detection algorithm [4, 7] is available to disaggregate the fingerprints or spectral envelope signatures of individual loads in the aggregate measurement.

An advantage of the NILM environment is the ability to monitor loads with minimal installation. To augment the capability of the NILM to include vibration and other environmental measurements, VAMPIRE provides *in-situ* monitoring for a load while powering itself from the magnetic fields around the wires energizing the load. A "mock-up" prototype of a VAMPIRE design is shown in Fig. 15. An inductively-coupled powered toroidal core with sensors is designed to collect and transmit information.

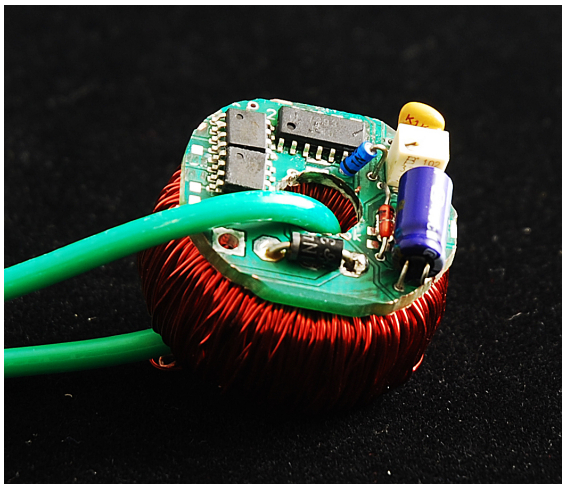

Fig. 15. A physical model of VAMPIRE. The current carrying wire running through the center of the core serves as a "primary".

The core can be physically mounted as part of the terminal box of the motor or other load, with one of the "hot" wires running through the center. This wire carrying I amps of current serves as the "primary" coil of this current transformer core. The "secondary" winding consists of N turns optimized to source sufficient power to the VAMPIRE monitor circuitry.

A circuit model of the current transformer is shown in Fig. 16 with the current I/N, leakage inductance L_ℓ, magnetizing inductance L_m, winding resistance R_w and load impedance Z_ℓ.

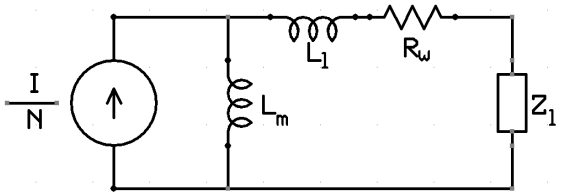

Fig. 16. Equivalent circuit model of the current transformer toroidal core. The current in the primary winding has been reflected over to secondary winding. The model includes a leakage inductance L_ℓ, a magnetizing inductance L_m, a winding resistance R_w, and a load impedance Z_ℓ.

For a fixed current I at a frequency $\omega = 2\pi f$ and secondary winding turns N, the power P_ℓ delivered to a load Z_ℓ is

$$P_\ell = \left| \frac{I}{N} \cdot \frac{j\omega L_m}{j\omega(Lm + L_\ell) + R_w + Z_\ell} \right|^2 Z_\ell. \quad (1)$$

To better understand the amount of power that can be harvested, consider a VAMPIRE around the wire energizing a 120 W load. With a 120 Volt-rms, $\omega = 2\pi \cdot 60$ radians utility line, the rms current drawn by the load is $I = 1$ A.

However, depending upon a magnitude of the current, a magnetic core might go into saturation, resulting in significantly less magnetic flux, which leads to less induced voltage and power for the VAMPIRE. Taking the core saturation effect into account, Fig. 17 shows the estimated amount of power that can be delivered to the secondary by the VAMPIRE. According to the graph, the available power is approximately 10 mW for a practical system.

The horizontal axis is Z_ℓ, the load impedance seen by the secondary winding of the core. The vertical axis is N, the number of secondary winding turns. The color gradient (third axis) represents the maximum power drawn in milliwatts, including core saturation.

For larger equipment, e.g., fractional or multi-horsepower motors, the power budget will be significantly greater. As the core is used to collect diagnostic information, a low power accelerometer, thermometer, microphone or any other desired sensor will be controlled by a low-power microcontroller (MCU) on the circular VAMPIRE pc-board. The MCU will collect all of the diagnostic information and report it.

Information could be collected wirelessly, or by a connection to the VAMPIRE board, or optically. In the NILM environment, a central monitor has access to all the power lines servicing the loads. If the information is inductively coupled onto one of the lines, this central unit can then receive the information as the data bits

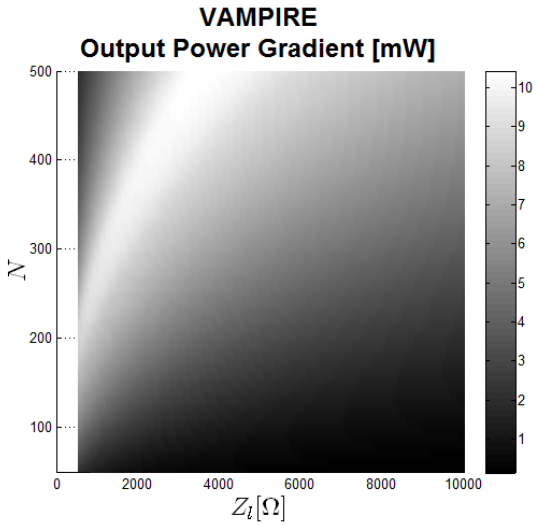

Fig. 17. This plot shows the power that can be inductively coupled from a 1 Amp, 120 W load. The vertical axis is N, the number of turns on the secondary winding and the horizontal axis is Z_ℓ, the load impedance connected to the secondary winding. The maximum VAMPIRE power is approximately 10 mW.

are sent over the power line. In all cases, the VAMPIRE design must provide adequate power to transmit data as desired.

VI. Future Work

Common methods for collecting data for system diagnostics are susceptible to error from a variety of sources. Hand sensors may be deployed incorrectly. Attempts to reduce data to figures-of-merit can hide important details. And individual "flavors" of sensors, e.g., vibration or temperature, may provide information that is best understood in the context of its correlation with other variables, e.g., power consumption.

Future hulls and retrofit improvements may benefit greatly from a "systems" approach to design. Sensors exemplified by the VAMPIRE concept presented here may be easy to incorporate in mission-critical devices. It is possible to create sensors that may exhibit improved reliability by requiring no special power or communication wiring. Information from these sensors can be correlated at a central point, and cross-referenced to reveal meaningful indicators that quickly expose pathologies. We also strongly believe that empowering crews to examine sensor data at scalable or adjustable levels of detail, as appropriate, to investigate a problem or casualty will result in improved fleet readiness with reduced operating costs.

VII. Acknowledgements

The authors grateful acknowledge the valuable advice and support of Deborah Nalchajian at ONR, The Grainger Foundation, Dr. Timothy J. McCoy, and the ESRDC program under Terry Ericsen at ONR. The acronym VAMPIRE arose from conversations at a workshop hosted by ONR. The authors are also grateful for the advice and support of USCG CAPT Paul Roden, CDR Edward Westfall, CDR John Gage, and Larry Wilkerson.

VIII. Appendix: Extracting Spectral Envelopes With The Hilbert Transform

The Hilbert transform is widely used in amplitude modulation systems for its ability to extract a narrowband signal that has been modulated by a carrier frequency. It is this ability that motivates the use of the Hilbert transform in measuring the acceleration levels shown in this paper.

The underlying assumption used throughout the paper is that there exists a narrowband of frequency content in the raw acceleration signal that can be used for diagnostic purposes. In Fig. 3, most of the acceleration from a ventilation fan in steady-state operation appears in those frequencies near the rotational speed of the fan.

After the raw acceleration signal is filtered by a 50-60 Hz passband filter, the spectral envelope of the acceleration shows the time evolution of the vibration levels from start-up then to steady-operation and finally to spin-down as shown in Fig. 4. It was also shown in Fig. 8 that analyzing the spectral content of the acceleration in the 30-40 Hz frequency band during the spin-down transient of the ventilation fan can be used to diagnose the underlying cause of any increase in vibration levels.

The Hilbert transforms makes it possible to compute these spectral envelopes and extract useful information from the rest of the acceleration signal.

An approach for computing the Hilbert transform, $\mathcal{H}_t\{\cdot\}$ can be seen in the frequency domain. Let the Fourier transform of $x(t)$ be $X(j\omega)$ and let the Fourier transform of the Hilbert transform, $\mathcal{H}_t\{x\}$ be $\mathcal{H}_\omega\{x\}$. The relationship between $x(t)$ and its Hilbert transform is as follows in the frequency domain: $\mathcal{H}_\omega\{x\} = (-j\mathrm{sgn}(\omega)) \cdot X(j\omega)$, where $\mathrm{sgn}(\cdot)$ is the sign function.

$$\mathcal{H}_\omega\{\cdot\} = \begin{cases} j, & \omega < 0 \\ 0, & \omega = 0 \\ -j, & \omega > 0. \end{cases} \tag{2}$$

In effect, the Hilbert transform applies a 90° phase shift to the negative frequency content and a -90° phase shift to the positive frequency content.

A simple example, shown in [8], gives a concrete example of how the Hilbert transform can be used to compute a spectral envelope. While it is assumed that useful content appears in a narrowband of frequency, for this example, consider the useful content exists at one frequency: ω_c. The amplitude (or vibration level) is the useful information and as it changes over time will be called $A(t)$ in the example.

Therefore, let $x(t) = A(t)\cos(\omega_c t)$ where the sinusoid of a carrier frequency ω_c is modulated by an amplitude envelope $A(t)$. Furthermore, assume $A(t)$ varies slowly compared to the "carrier" frequency and can therefore be approximated as some constant for simplicity. The Hilbert transform for $x(t)$ as shown in [8] is

$$y(t) \approx A(t)\sin(\omega_c t). \tag{3}$$

The next step to obtaining the envelope in amplitude modulation scenarios involves computing a signal $z(t)$ as such:

$$z(t) = x(t) + j\mathcal{H}_t\{x\}. \tag{4}$$

The signal $z(t)$ is called an analytic signal and further details can be found in [8]. In the example,

$$z(t) \approx A(t)(\cos(\omega_c t) + j\sin(\omega_c t)) \tag{5}$$

$$\approx A(t)e^{j\omega_c t}, \tag{6}$$

where the last equality holds from Euler's relation. Once $z(t)$ is computed, obtaining the envelope $A(t)$ is found by taking the absolute value as

$$|z(t)| \approx \left|A(t)e^{j\omega_c t}\right| \tag{7}$$

$$\approx |A(t)|\,. \tag{8}$$

The resultant signal $|A(t)|$ is the envelope of the underlying narrowband signal which is used to track vibration levels in this paper.

REFERENCES

[1] E. Linton, "Coast Guard RCM/CBM Initiatives." Presented at the 2005 Department of Defense Maintenance Symposium and Exhibition, Birmingham, AL, Oct. 2005.

[2] W. T. Thomson, R. A. Leonard, A. J. Milne, and J. Penman, "Failure Identification of Offshore Induction Motor Systems Using On-condition Monitoring," *Reliability Engineering*, vol. 9, no. 1, pp. 49–64, 1984.

[3] M. DiUlio, C. Savage, B. Finley, and E. Schneider, "Taking the Integrated Condition Assessment System to the Year 2010," in *13th Int. Ship Control Systems Symposium*, Orlando, FL., 2003.

[4] S. B. Leeb, S. R. Shaw, and J. L. Kirtley Jr., "Transient Event Detection in Spectral Envelope Estimates For Nonintrusive Load Monitoring," *IEEE Trans. Power Del.*, vol. 10, no. 3, pp. 1200–1210, Jul 1995.

[5] S. B. Leeb, "A Conjoint Pattern Recognition Approach to Nonintrusive Load Monitoring," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, Feb. 1993.

[6] S. R. Shaw, "System identification techniques and modeling for nonintrusive load diagnostics," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, Feb. 2000.

[7] R. Cox, S. B. Leeb, S. R. Shaw, and L. K. Norford, "Transient Event Detection For Nonintrusive Load Monitoring and Demand Side Management Using Voltage Distortion," Mar. 2006.

[8] J. O. Smith III, *Mathematics of the Discrete Fourier Transform (DFT): with Audio Applications – Second Edition*. W3K Publishing, 2007.