

Multistream Power Monitoring for Energy Systems

Aaron W. Langham, Thomas C. Krause, Daisy H. Green, Steven B. Leeb

Abstract—Power monitoring provides services such as billing, power system diagnostics, condition-based maintenance, and continuous commissioning. Many monitoring applications assume that grid voltage conditions are “stiff” while analysis focuses on load variations. Future power networks will more prominently include microgrids and distributed energy resources. Although power monitoring will be more important to these systems, it will also be made more difficult by lower generation inertia, higher source impedance, and increased exposure to generation control methods. This work provides a re-examination of a high bandwidth power monitor’s preprocessing front-end for systems without a stiff voltage waveform. Preprocessing techniques are presented for adapting to nonstationary voltage amplitudes and frequencies. Simultaneous use of these techniques creates a “multistream” strategy to power monitoring, in which several data streams are produced and the most appropriate stream(s) for a given monitoring task can be selected.

Index Terms—Power monitoring, low-inertia power systems, microgrids, data preprocessing

I. INTRODUCTION

Microgrids, distributed energy resources (DERs), and conventional bulk power systems benefit from power monitoring [1], [2]. Monitoring enables usage tracking, billing, and condition-based maintenance [3], [4]. Machine learning and automated systems can supplement or replace human operators in performing these tasks [5]. Both machine learning systems and human operators perform best with properly presented data. Nonstationary data distributions drastically complicate both training and inference for a machine learning system. Preprocessing techniques can mitigate or exacerbate this issue. Preprocessing includes a multitude of data manipulation techniques including filtering, basis transformation, and resampling [6]. In the context of ac electrical energy systems and microgrids in particular, currents and voltages are typically sampled at the kilohertz level. However, this data is highly redundant, as most electrical signals of interest effectively “modulate” a 60 Hz carrier wave and its harmonics [7]. Preprocessing is important for power monitoring because it can exploit these physics-based properties to compress the data to a lower dimension without losing meaningful information.

In electrical systems, the distribution of power data may vary due to changing grid conditions. Source-driven variations include generation voltage waveform deviations, both in amplitude and frequency. Load-driven variations include dynamic power consumption and harmonic content. For a given monitoring problem, only one of these variations may actually be of interest. Nonintrusive load monitoring serves as an illustrative example. A nonintrusive load monitor (NILM) disaggregates a system’s power consumption profile into individual load activity. A wealth of techniques have been developed for load identification in residential settings [8]. NILM techniques have

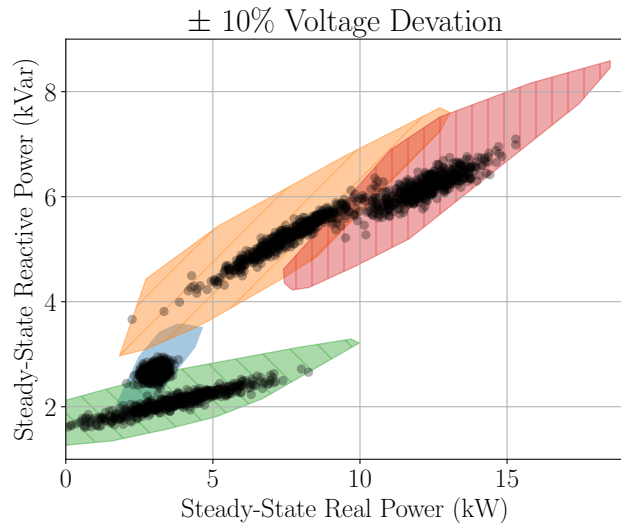


Fig. 1: Visualization of source-driven variations in voltage amplitude that lower cluster separability.

also been developed for load forecasting [9] and load scheduling [10]. In industrial settings, NILMs have been shown to be useful for performing fault detection and diagnostics on loads [11]. These applications are primarily concerned with load-driven variations, yet are confounded by source-driven variations. To illustrate, Fig. 1 shows an example visualization of source-driven variations in voltage amplitude. This data was generated to match the distributions of four shipboard electrical loads presented in [12], modeled to be linear. Each black dot represents an “ON” electrical event, plotted as the steady-state difference in real and reactive power. The shaded polygons show the convex hulls of the event distributions after applying a $\pm 10\%$ deviation in the system voltage amplitude. With this added ambiguity, the individual clusters of loads overlap and become more difficult to separate. This presents a problem for identification and diagnostic techniques that rely on separating data for different loads. For example, for diagnostics it is important to be able to distinguish source-driven changes in load power from changes due to degradation or other dynamic conditions, such as changes in system operating pressures or temperatures.

This work addresses source-driven variations due to deviations in the frequency and amplitude of system voltage waveforms. The contributions of this paper are as follows:

- 1) Techniques for adapting conventional spectral envelope preprocessing to grids with nonstationary frequencies and voltage amplitudes.
- 2) Mathematical characterizations of these techniques’ utility across different types of electrical loads.

- 3) An evaluation of these techniques on data collected on a shipboard microgrid.

Section II gives a review of power computation techniques and spectral envelope processing. Section III discusses design choices in adapting spectral envelope processing to nonstationary grid frequencies. Section IV develops preprocessing techniques that correct for voltage amplitude variations, which can be combined into a preprocessing system outputting multiple data streams. Section V gives an examination of these techniques' suitability for different loads. Finally, Section VI presents data from a marine microgrid to demonstrate the efficacy of these techniques.

II. POWER COMPUTATION METHODS

Many power computation techniques exist to extract useful signals and figures of merit for ac electrical energy systems [13]. An elementary measure of power is instantaneous power transfer, which is simply the product of voltage and current. A rolling average of instantaneous power yields *average* real (or active) power, useful for billing purposes and estimating long-term physical work done by a system. The product of the root-mean-square (rms) values of voltage and current yields *apparent* power, which is useful for sizing generation and transmission components. A lag or lead in the current waveform with respect to the voltage waveform produces a displacement factor, representing energy storage in the system. Currents at frequencies other than the utility frequency produce a distortion factor, indicating harmonic content. Although some fields use the term reactive power to refer to both [14], this work reserves the term for periodic energy storage at the fundamental frequency due to a displacement factor. Phasor quantities for real power (P), reactive power (Q), and apparent power (S) simplify analysis for sinusoidal ac grids. Classical power system analysis tracks these quantities in steady-state grid operation.

Modern power monitoring techniques for ac grids extend these notions to dynamic grid behavior, such as changes in plant or load profile and random disturbances [15]. These dynamics often occur over a few utility line cycles. As a result, they are effectively invisible to steady-state analysis. Power spectral envelopes capture fine-grained temporal behavior and enhance the resolution of electrical features from current measurements [7], [16]. When appropriately scaled, the fundamental quadrature and in-phase spectral envelopes approximate real and reactive power in steady state. Higher-order harmonic spectral envelopes preserve information about current waveform distortions. For a system with a sinusoidal voltage waveform, average power flows only at the fundamental frequency. However, scaling harmonic current spectral envelopes by a voltage scale factor produces values with units of power, dimensionally compatible with the more physically related fundamental quadrature and in-phase spectral envelopes. For both the fundamental and higher-harmonic cases, the choice of voltage scale factor is a design choice that can target different monitoring goals.

Consider one window of time corresponding to a line cycle of an ac system's voltage waveform, starting at the positive

zero crossing. Sampling the current waveform over this time yields sequence $i[n]$ with number of sampling points N . The following equations produce spectral envelopes P_k , Q_k , at harmonic order k :

$$P_k = -\frac{V_s}{N} \text{imag}\{I(kf_0)\}, \quad (1)$$

$$Q_k = \frac{V_s}{N} \text{real}\{I(kf_0)\}. \quad (2)$$

Here, f_0 corresponds to the grid frequency and $I(kf_0)$ corresponds to the discrete-time Fourier transform (DTFT) of $i[n]$, sampled at kf_0 . The voltage scale factor V_s transforms units of current into units of power. For $k = 1$ and V_s equal to the system voltage amplitude (denoted V_{pk}), P_1 and Q_1 approximate the traditional phasor values of real and reactive power in steady state. This system is *phase-aware* and requires sensing both the voltage and current. A system that only measures current is limited to computing apparent power spectral envelopes S_k :

$$S_k = \frac{V_s}{N} |I(kf_0)| = \sqrt{P_k^2 + Q_k^2}. \quad (3)$$

This type of computation is *phase-unaware*. Variables V_{pk} and f_0 are nearly constant on systems with stiff voltage waveforms. However, these quantities deviate from their nominal values on all practical grids, especially microgrids. These deviations may be constant or time-varying in nature. Although V_{pk} and f_0 may contain deviations faster than the grid frequency, this work only considers the case in which V_{pk} and f_0 are locally constant for the duration of one line cycle.

This work presents five power spectral envelope processing techniques that adapt to non-ideal voltage waveforms. The first technique adapts to variations in system frequency (f_0). The other four techniques compensate for variations in voltage amplitude (V_{pk}). Each technique is useful for a different purpose, including enhancing the success of machine learning recognition and condition-based maintenance. In many systems, abundant computation power allows a designer to "mix and match" these techniques and produce multiple data streams, referred to here as a "multistream" approach. The most useful streams for a given monitoring task can be retrieved in later analysis. The frequency adaptation and amplitude adaptation techniques are independent. Thus, both frequency and amplitude adaptation can be achieved by using the technique in Section III in conjunction with one of the techniques in Section IV.

III. FREQUENCY ADAPTATION

Power spectral envelope computation requires an estimate of the system frequency to properly sample the DTFT. This frequency estimate \hat{f}_0 is typically either obtained in real time with sensors or approximated as the nominal value (e.g., 60 Hz in the US). Differences between the estimated and actual system frequency introduce artifacts in the computed power spectral envelopes. The magnitude of the frequency estimation error dictates the size and severity of these artifacts. Depending on the system requirements, these artifacts may warrant extra sensing to obtain a better frequency estimation. This section




Fig. 2: Oscillatory artifacts in phase-unaware spectral envelopes as a result of incorrect system frequency estimation.

characterizes the size and character of the artifacts introduced by deviations in grid frequency for both phase-unaware and phase-aware spectral envelope computation.

A. Phase-Unaware

A phase-unaware monitoring scheme only measures currents. This facilitates a minimal sensor profile and lowers hardware requirements. A phase-unaware system computes apparent power spectral envelopes (S_k) from the DTFT of $i[n]$ as shown in Eq. (3). Although this method loses phase information such as power factor, it preserves magnitude estimates of fundamental and harmonic currents. If the actual and estimated grid frequencies do not match, the output S_k stream contains an oscillatory artifact. Fig. 2 shows this artifact for a simulated 1000 VA load on a nominally 60 Hz system, as the system frequency steps from 59 Hz to 61 Hz.

The frequency of this artifact is approximately equal to twice the difference in expected and actual system frequencies. To explain this analytically, first assume that a continuous-time integral producing the function $S(t)$ approximates the output stream of (3) for $k = 1$:

$$S(t) = \frac{V_s}{\hat{T}_0} \left| \int_t^{t+\hat{T}_0} i(\tau) e^{-j2\pi\hat{f}_0\tau} d\tau \right|, \quad (4)$$

where $\hat{T}_0 = 1/\hat{f}_0$ and \hat{f}_0 is the estimated system frequency, $i(\tau)$ is characterized as $I_{pk} \cos(2\pi f_0\tau)$, and f_0 is the actual system frequency. The actual current waveform may have a different phase; this analysis takes it as a cosine wave without loss of generality. This results in the following integral:

$$S(t) = \frac{V_s I_{pk}}{\hat{T}_0} \left| \int_t^{t+\hat{T}_0} \cos(2\pi f_0\tau) e^{-j2\pi\hat{f}_0\tau} d\tau \right|. \quad (5)$$

Transforming the cosine term into a sum of conjugate exponentials and performing the integration yields the following:

$$S(t) = \frac{V_s I_{pk}}{4\pi\hat{T}_0} \left| \left[\frac{e^{j2\pi(f_0-\hat{f}_0)\tau}}{f_0 - \hat{f}_0} - \frac{e^{-j2\pi(f_0+\hat{f}_0)\tau}}{f_0 + \hat{f}_0} \right]_{t}^{t+\hat{T}_0} \right|. \quad (6)$$

Carrying out this integral and using the binomial approximation when computing the magnitude results in an expression for $S(t)$ in the following form:

$$S(t) \approx S_0 \left(1 + \frac{\hat{f}_0^2 - f_0^2}{2(\hat{f}_0^2 + f_0^2)} \cos(4\pi(\hat{f}_0 - f_0)t + \phi) \right), \quad (7)$$

where S_0 is the correct apparent power value and ϕ is a phase term unimportant to this analysis. From this expression, it is clear that the measured apparent power spectral envelopes oscillate around S_0 at a frequency equal to twice the difference between the estimated and actual system frequencies. In addition, the amplitude of the artifact as a fraction of S_0 is given by $(\hat{f}_0^2 - f_0^2)/(2(\hat{f}_0^2 + f_0^2))$. For systems with constraints on expense and computation, a phase-unaware spectral envelope processing system may be entirely sufficient or may introduce prohibitively large artifacts. A designer can use this analysis to determine whether a phase-unaware system is appropriate for a given monitoring application.

B. Phase-Aware

A phase-aware monitoring scheme requires both voltage and current sensors. This scheme computes real (P_k) and reactive (Q_k) power spectral envelopes using Eqs. (1) and (2). Just as in the phase-unaware case, this scheme requires an estimate of the grid frequency, and errors in this estimate introduce artifacts in the output P_k and Q_k streams. To illustrate these artifacts, Fig. 3 shows the results of a series of step changes in system frequency from 59 Hz to 61 Hz for a simulated 1000 W resistive load on a nominally 60 Hz system. Although not sinusoidal, these artifacts introduce uncertainty into measurements of real and reactive power. This uncertainty complicates further efforts at load identification and diagnostics.

To derive the size of these artifacts as a function of frequency estimation error, consider the spectral envelope computation window shown in Fig. 4. This plot shows current waveforms for four loads with different lag angles denoted by θ . The following integral for fundamental complex power ($\bar{S}_1 = P_1 + jQ_1$) provides an approximation of phase-aware spectral envelopes for this window:

$$\bar{S}_1 = \frac{V_s}{j\hat{T}_0} \int_0^{\hat{T}_0} i(t) e^{j2\pi\hat{f}_0 t} dt. \quad (8)$$

Substituting $I_{pk} \sin(2\pi f_0 t - \theta)$ as $i(t)$ yields:

$$\bar{S}_1 = \frac{V_s I_{pk}}{j\hat{T}_0} \int_0^{\hat{T}_0} \sin(2\pi f_0 t - \theta) e^{j2\pi\hat{f}_0 t} dt \quad (9)$$

Performing this integral results in the following expression for P_1 and Q_1 :

$$P_1 = \frac{V_s I_{pk} \hat{f}_0^2}{\pi(\hat{f}_0^2 - f_0^2)} \sin\left(\pi \frac{f_0}{\hat{f}_0}\right) \cos\left(\pi \frac{f_0}{\hat{f}_0} - \theta\right), \quad (10)$$




Fig. 3: Constant artifacts in phase-aware spectral envelopes as a result of inaccurate system frequency estimation.




Fig. 4: Example spectral envelope computation window when the expected system frequency differs from the actual system frequency, for multiple load power factors. The black line at time 1.0 represents the actual period of the current waveform, which in this figure is shorter than the expected period.

$$Q_1 = -\frac{V_s I_{pk} f_0 \hat{f}_0}{\pi(f_0^2 - \hat{f}_0^2)} \sin\left(\pi \frac{f_0}{\hat{f}_0}\right) \sin\left(\pi \frac{f_0}{\hat{f}_0} - \theta\right). \quad (11)$$

By altering the values of P_1 and Q_1 , these artifacts distort the power factor of the system load. The sensitivity of the output P_1 and Q_1 with respect to the frequency estimation




Fig. 5: Plots of P_1 and Q_1 versus f_0 for $\hat{f}_0 = 60$ Hz and $\theta = 0$, as given in Eqs. (10) and (11). For this load and system, the Q_1 curve is steeper and more sensitive to deviations in f_0 . The black dots show the operating points shown in Fig. 3.

error depends on the load displacement angle θ . For a resistive load (i.e. $\theta = 0$), Q_1 is more sensitive to errors in frequency estimation than P_1 , which can be seen in Fig. 3. This is due to the expressions for P_1 and Q_1 having different sensitivities to changes in f_0 , as shown in Fig. 5. If the grid frequency is known to exist in a given range, a preprocessing mechanism could mitigate these deviations by using Eqs. (10) and (11) to find the maximum amount of deviation for both P_1 and Q_1 for the system, and disregarding electrical transients below this value.

In some systems, the artifacts introduced in the previous sections may be entirely tolerable. Otherwise, the power computation system must estimate the instantaneous grid frequency to properly sample the DTFT. Intuitively, if the grid transmits power on a given frequency that shifts in time, the power computation system should adapt as well. Fig. 6 shows the phase-aware monitoring scenario of Fig. 3, except with a real-time frequency estimate available to the power computation scheme. Although the system frequency steps from 59 Hz to 61 Hz, the outputs are resilient to artifacts (except at the points at which the frequency steps).

IV. AMPLITUDE ADAPTATION

A great volume of research on nonintrusive load monitoring neglects the effect of voltage amplitude variability [17]. For high-inertia, low-impedance power systems such as a bulk power system, the assumption that the voltage amplitude is “stiff” is often accurate enough to be useful. However, microgrids have lower inertia than bulk power systems, since smaller machines or power electronics-controlled resources usually provide their generation. Accordingly, the voltage amplitude is less tightly controlled in microgrids than in bulk power systems. Microgrids also typically have a higher equivalent source impedance than the bulk power system since there are fewer generation sources. As loads draw current, the supply voltage drops across the source impedance before




Fig. 6: Resulting real and reactive power from an adaptive power computation scheme while the system frequency steps from 59 Hz to 61 Hz.

appearing across the load. Source-driven variations in voltage amplitude often result in a corresponding variation in current consumption. The choice of the V_s scale factor in Eqs. (1), (2), and (3) determines how “sensitive” the power computation’s output is to these changes. This section explores three different choices for the scale factor V_s : the nominal value, the measured value, and a scaled value that nulls amplitude deviations. In addition, this section presents a technique for loads with constant power components. This section characterizes these four techniques mathematically. Section V provides further insight on the utility of these techniques for different loads. These techniques come secondary to those of Section III and require an estimate of the system frequency, either from sensing or a stiff approximation. All four (or some subset) of these techniques can be computed, producing multiple output data streams.

A. The Nominal Spectral Envelope

The simplest choice of V_s is simply the nominal voltage amplitude, e.g. $120V\sqrt{2} \approx 170$ V on US residential grids. This work refers to this choice as producing a *nominal spectral envelope*. Many shipboard NILM deployments use this scheme [18] and forego the expense associated with voltage amplitude estimation. However, this method produces spectral envelope values that do not precisely match the physical quantities of real and reactive power consumption when the voltage amplitude differs from its nominal value. Fig. 7 shows deviations introduced into the output real power stream as the voltage

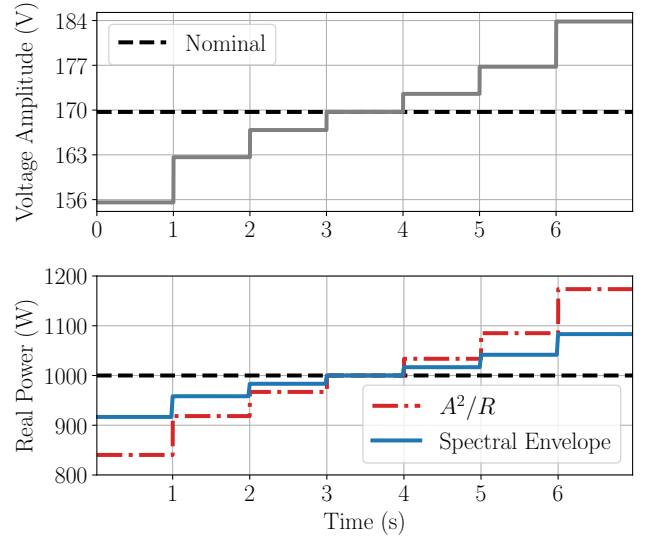


Fig. 7: Artifacts introduced in simulated non-adaptive power computation from step changes in the voltage amplitude.

amplitude (denoted A) is stepped on a nominally 120 V (rms) system with a resistive 1000 W load. As expected of a linear system the power consumption changes as the square of the system voltage. However, the changes in the spectral envelope stream do not completely match the expected changes in physical real power consumption (i.e., $A^2/(2R)$).

To examine this discrepancy, consider a phasor model of the power system. An impedance in the phasor domain with $\bar{Z} = |Z|\angle\theta_Z$ characterizes the system load. The phasor $\bar{V} = A/\sqrt{2}\angle 0$ represents the system voltage, where A is the system’s actual, possibly deviating, voltage amplitude. The phasor domain current is therefore:

$$\bar{I} = \frac{\bar{V}}{\bar{Z}} = \frac{A}{|Z|\sqrt{2}}\angle -\theta_Z. \quad (12)$$

Nominal spectral envelope computation for the fundamental frequency approximates the following:

$$\bar{S}_1 = \frac{V_N}{\sqrt{2}}\bar{I}^* = A\frac{V_N}{2|Z|}\angle\theta_Z, \quad (13)$$

where V_N is the system’s nominal voltage amplitude. Real and reactive power spectral envelopes are:

$$P_1 = \text{real}\{\bar{S}_1\} = A\frac{V_N}{2|Z|}\cos\theta_Z, \quad (14)$$

$$Q_1 = \text{imag}\{\bar{S}_1\} = A\frac{V_N}{2|Z|}\sin\theta_Z. \quad (15)$$

This stands in contrast to the actual physical real (P) and reactive (Q) power drawn by the load, which are proportional to A^2 rather than A . Therefore, for a linear system, the P_1 and Q_1 of this technique differ from the physical values of P and Q by a factor of V_N/A (that is, $P_1 = PV_N/A$).

B. The Sensed Spectral Envelope

To produce spectral envelope values P_1 and Q_1 that are consistent with physical real and reactive power consumption, this next approach measures the voltage amplitude for each spectral envelope computation using voltage sensors, producing a *sensed spectral envelope*. Representing the voltage amplitude again as A , this technique sets the scale factor V_s to be A for each spectral envelope computation. Using the same phasor analysis and linear system as in Section IV-A, the value of \bar{S}_1 will approximate the following:

$$\bar{S}_1 = \frac{A}{\sqrt{2}} \bar{I}^* = \frac{A^2}{2|Z|} \angle \theta_Z. \quad (16)$$

Dividing this into real and reactive power:

$$P_1 = \text{real}\{\bar{S}_1\} = A^2 \frac{\cos \theta_Z}{2|Z|}, \quad (17)$$

$$Q_1 = \text{imag}\{\bar{S}_1\} = A^2 \frac{\sin \theta_Z}{2|Z|}. \quad (18)$$

Spectral envelopes that match physical quantities are useful for certain applications, such as power metering for billing. However, consider a power monitoring system designed to perform load identification and diagnostics. Whenever the system voltage amplitude increases, either suddenly or over time, the power consumed by a linear load increases quadratically. This is apparent in the equations for the P_1 and Q_1 of this technique – the spectral envelope values are quadratically sensitive to A . This is potentially disastrous for load identification and diagnostics, as a power monitor will perceive this load as a higher-power load, despite its identity and condition being unchanged. This frustrates a NILM's load identification capability, since faults are sometimes found by examining subtle changes in steady-state or transient power consumption in spectral envelopes.

C. The Compressed Spectral Envelope

The third scaling technique aims to undo or compress the effect of system voltage amplitude variations, creating a *compressed spectral envelope*. To do so, it estimates the aggregate load impedance \bar{Z} over one line cycle using measured voltage amplitude A and phasor current \bar{I} :

$$\bar{Z} = \frac{\bar{V}}{\bar{I}} = \frac{A}{\bar{I}\sqrt{2}}. \quad (19)$$

Using \bar{Z} , this technique estimates the complex power that would be consumed by this load at the nominal voltage amplitude V_N :

$$\bar{S}_1 = \frac{V_N^2}{2\bar{Z}^*} = \left(\frac{V_N^2}{A}\right) \frac{\bar{I}^*}{\sqrt{2}} = \frac{V_N^2}{2|Z|} \angle \theta_Z. \quad (20)$$

Matching this equation to Eqs. (13) and (16) and their respective values of V_s shows that this is equivalent to using $V_s = V_N^2/A$. This produces the following P_1 and Q_1 :

$$P_1 = \text{real}\{\bar{S}_1\} = \frac{V_N^2}{2|Z|} \cos \theta_Z, \quad (21)$$

$$Q_1 = \text{imag}\{\bar{S}_1\} = \frac{V_N^2}{2|Z|} \sin \theta_Z. \quad (22)$$

These values of P_1 and Q_1 are no longer functions of A , “compressing” the effect of voltage amplitude variations. Importantly, this technique is not ideal for energy scorekeeping or estimation, since the spectral envelope values it produces are no longer designed to match physical power consumption. However, for load identification and diagnostics, this technique provides a set of load fingerprints that are independent of exogenous voltage amplitude deviations.

D. The Affine Spectral Envelope

The previous techniques assume a linear view of a grid's load. This assumption may fail for loads designed to draw an amount of power that is independent of the voltage amplitude, such as power electronics. In practice, they typically also draw real and reactive power proportional to the square of the voltage. This work models the power consumption of these loads as an affine function of A^2 :

$$P = \alpha_P + \beta_P A^2, \quad (23)$$

$$Q = \alpha_Q + \beta_Q A^2. \quad (24)$$

In this model, α_P and α_Q represent the constant power consumption. Variables β_P and β_Q represent the power consumption component that behaves as a “linear load.” For convenience in later analysis, let γ_P and γ_Q be defined as the ratio of constant power consumption to linear power consumption at the nominal voltage amplitude, for both real and reactive power:

$$\gamma_P = \frac{\alpha_P}{\beta_P V_N^2}, \quad \gamma_Q = \frac{\alpha_Q}{\beta_Q V_N^2}. \quad (25)$$

For loads with an entirely constant real power consumption (i.e. $\beta_P = 0$), $\gamma_P = \infty$. Linear loads, for which $\alpha_P = \alpha_Q = 0$, correspond to $\gamma_P = \gamma_Q = 0$. The final preprocessing approach, referred to as an *affine spectral envelope*, removes a known constant power offset for a given “affine” load, computes the compressed spectral envelopes of Section IV-C, and then adds the constant power offset. In this way, only the “linear” component of the load is rescaled. Let P_c and Q_c represent the compressed spectral envelopes computed over a given line cycle. The equations for the affine spectral envelope are:

$$P_1 = P_c + \alpha_P \left(1 - \frac{V_N^2}{A^2}\right). \quad (26)$$

$$Q_1 = Q_c + \alpha_Q \left(1 - \frac{V_N^2}{A^2}\right). \quad (27)$$

As with the compressed spectral envelope, these values of P_1 and Q_1 are (by design) no longer accurate estimates of real and reactive power. Instead, they model the load's expected real and reactive power signature at nominal voltage amplitude conditions. In a sense, this technique is a generalization of compressed spectral envelopes, allowing for non-zero values of α_P and α_Q .

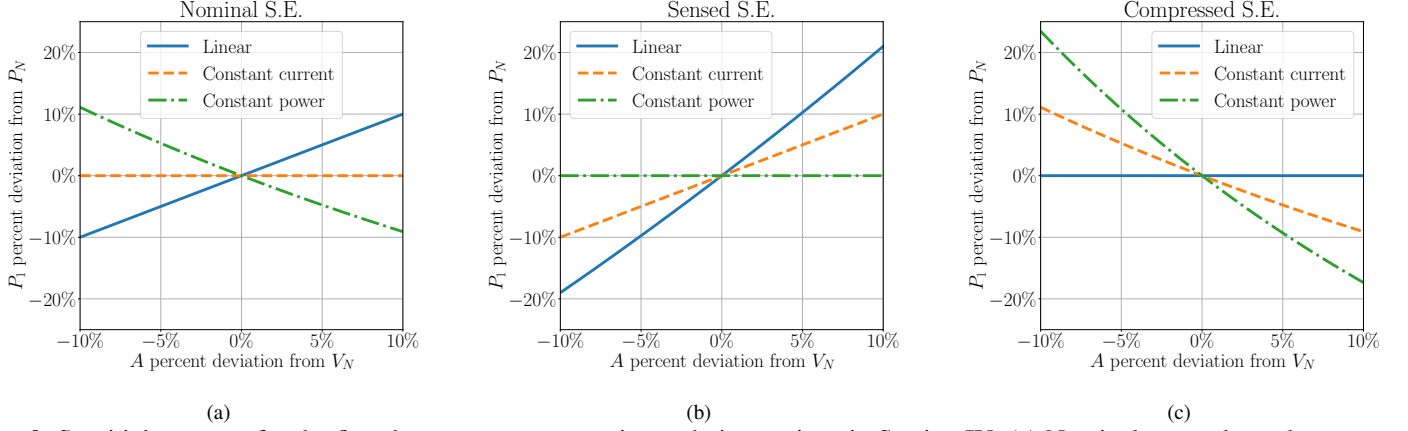


Fig. 8: Sensitivity curves for the first three power computation techniques given in Section IV. (a) Nominal spectral envelopes. (b) Sensed spectral envelopes. (c) Compressed spectral envelopes.

Load type	Preprocessor Technique			
	Nominal	Sensed	Compressed	Affine
Constant current	0	1	1	-
Constant power	1	0	2	0
Linear	1	2	0	0
Affine	$\frac{ \gamma-1 }{ \gamma+1 }$	$\frac{2}{ \gamma+1 }$	$\frac{2 \gamma }{ \gamma+1 }$	0

TABLE I: Sensitivity for each voltage amplitude scaling technique across constant current, constant power, linear, and affine loads.

V. SENSITIVITY ANALYSIS

These techniques' sensitivities to actual voltage amplitude A reveal which are most useful for a given monitoring application. The following sensitivity analysis considers real power spectral envelopes without loss of generality. However, the analysis also applies for reactive and apparent spectral envelopes. For real power spectral envelope P_1 , let P_N denote its value at a nominal voltage amplitude. The percent difference of P_1 from P_N is $(P_1 - P_N)/P_N$. Similarly, for measured voltage amplitude A and nominal voltage amplitude V_N , the percent difference is $(A - V_N)/V_N$. The derivative of the percent difference in P_1 with respect to the percent difference in A yields the sensitivity, as follows:

$$\frac{\partial(P_1 - P_N)/P_N}{\partial(A - V_N)/V_N} = \frac{V_N}{P_N} \frac{\partial P_1}{\partial A}. \quad (28)$$

Finally, taking the absolute value and evaluating at $A = V_N$ yields:

$$\text{Sensitivity} = V_N \left| \frac{1}{P_N} \frac{\partial P_1}{\partial A} \right|_{A=V_N}. \quad (29)$$

Linear loads obey Ohm's law and are characterized by an impedance $|Z| \angle \theta_Z$. For these loads, $P_N = V_N^2 \cos \theta_Z / (2|Z|)$. Making this substitution and evaluating at $A = V_N$ yields the following expression for sensitivity:

$$\text{Sensitivity} = \frac{2|Z|}{V_N} \left| \frac{1}{\cos \theta_Z} \frac{\partial P_1}{\partial A} \right|_{A=V_N}. \quad (30)$$

Substituting for P_1 for each technique yields the sensitivity of each technique for linear loads. Table I shows the results

of similar analysis for constant current amplitude, constant power, linear, and affine loads. Quadratic sensitivity, e.g. for linear loads using the sensed spectral envelope, corresponds to a sensitivity value of 2. The sensitivity of nominal, sensed, and compressed spectral envelopes for affine loads is a function of γ_P . For constant current loads, there are no well-defined values of α and β , so affine spectral envelopes cannot be applied.

Notably, six combinations result in an output P_1 that is insensitive to A , as shown in bold in Table I. To illuminate this, Fig. 8 shows simulated sweeps of voltage amplitude and the corresponding preprocessor output, for each combination of load type and technique. Affine loads and affine spectral envelopes are not shown, since affine loads' sensitivity depends on γ_P and affine spectral envelopes are not well-defined for constant current loads. In Fig. 8a, the output P_1 is insensitive to amplitude deviations for constant current loads, and approximately linearly sensitive to amplitude deviations for linear and constant power loads. Fig. 8b shows a similar insensitivity for constant power loads, but an especially high (approximately quadratic) sensitivity for linear loads. Finally, Fig. 8c shows that the final technique is insensitive for linear loads, but approximately quadratically sensitive for constant power loads. In a sense, the technique of Section IV-A provides a trade-off, where linear and constant power loads will always be sensitive to deviations in voltage amplitude, but never quadratically. A designer can choose one of these spectral envelope processors based on the load composition or compute all four in real time and select the most useful one in later analysis.

VI. DEMONSTRATION

This section presents actual field data from a shipboard microgrid to illustrate the utility of the presented amplitude scaling techniques. The proposed processing addresses characteristics of nearly all microgrids. For example, microgrids typically switch between "grid-connected" and "islanded" operating modes [19]. The character and statistical distribution of the voltage supply waveform will likely change discretely when the operating mode is changed. This phenomenon is

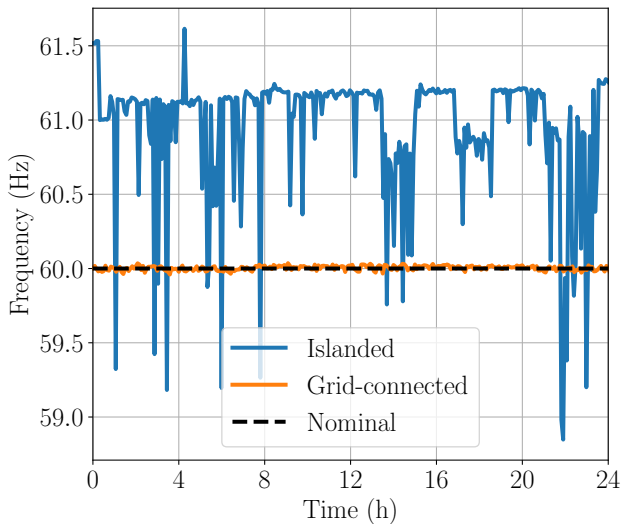


Fig. 9: Measured system frequency on USCGC Sturgeon over a one-day period for both islanded and grid-connected configurations. One frequency sample is shown every five minutes.

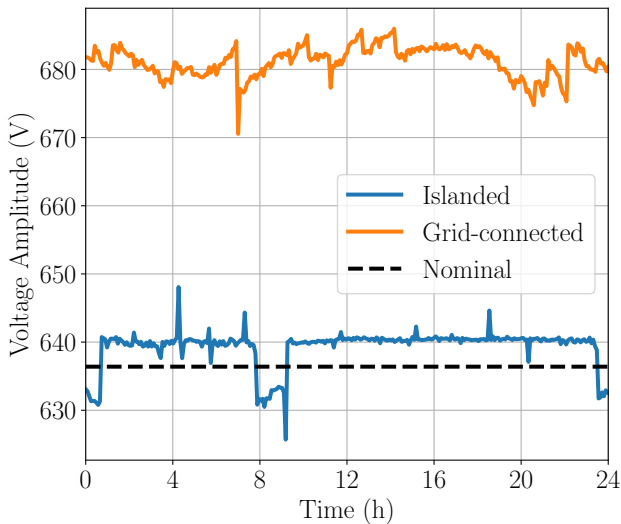


Fig. 10: Measured system voltage amplitude on USCGC Sturgeon over a one-day period for both islanded and grid-connected configurations. One voltage sample is shown every five minutes.

exhibited in shipboard electrical systems as they switch generation sources when getting underway or coming in port [20]. In addition, microgrids often contain lower-inertia generation sources that make voltage waveform regulation difficult. Ship power systems are inherently low-inertia when underway, as they typically rely on a select set of small generators in their plant lineup [21], [22]. On shipboard grids, the system voltage waveform's character varies constantly due to changing load, generator paralleling, and switching operating modes. For instance, load dynamics cause large changes in voltage amplitude and frequency as generators and their controls cannot

instantly respond to changes in demand. Finally, droop control allows multiple generators to share the ship's electrical load, but by design the system frequency drops as the load increases [23].

Fig. 9 shows a comparison of grid frequency quality across microgrid operation modes for US Coast Guard Cutter (USCGC) Sturgeon, an 87 ft. (27 m.) patrol boat. This plot shows frequency data for 24-hour periods of both islanded (underway) and grid-connected (in port) operation. The dashed black line indicates the nominal 60 Hz value. The frequency is clearly better regulated in grid-connected mode than in islanded mode. This ship uses two three-phase ac generators, which provide much less inertia than the terrestrial grid. Fig. 10 shows an analogous comparison of voltage amplitudes across modes. Again, the islanded and grid-connected values differ significantly. In islanded mode the plant creates a nominally 450 V grid (measured line-to-line, rms). When on shore power (i.e., grid-connected) the terrestrial bulk power system supplies the ship with 480 V (line-to-line, rms). The average voltage amplitude differs between the two modes by approximately 42 V, which is more than 6% of the nominal 636 V value. The adaptation techniques of Section III and Section IV are demonstrated in this section in order to compensate for these inconsistent voltage waveforms.

A. Separability Analysis

This section examines electrical events identified on USCGC Sturgeon across a 15-day period, in which the ship operated in both underway and in-port modes. Fig. 11 shows steady-state changes in real and reactive power of each electrical event using the first three techniques of Section IV. The three plots all show the same physical events, but using different preprocessing. These events correspond to five distinct shipboard loads with unique real and reactive power consumption profiles. Thus, all three preprocessing methods are able to distinguish these events into distinct clusters. However, the distance between clusters differs in each plot, suggesting that classification difficulty varies across all three techniques. This distance between clusters, referred to as *separability*, is important to the ability of data-driven techniques to support load identification and diagnostics. Since load electrical behavior can drift over time, clusters of loads may become close enough to be indistinguishable [12]. As shown in Section I, voltage amplitude variations cause electrical events to occupy an enlarged feature space region. The silhouette score is a figure of merit for cluster separability [24]. The following equation computes the silhouette score s across D data points:

$$s = \frac{1}{D} \sum_{i=1}^D \frac{b_i - a_i}{\max(a_i, b_i)}, \quad (31)$$

where a_i is the average Euclidean distance between point i and the other data points in the same cluster, and b_i is the average distance between point i and each point in the next-closest cluster. This score is bounded by $[-1, 1]$, with a higher score representing a highly separable set of clusters. The silhouette scores for Fig. 11a, Fig. 11b, and Fig. 11c are 0.60, 0.66, and

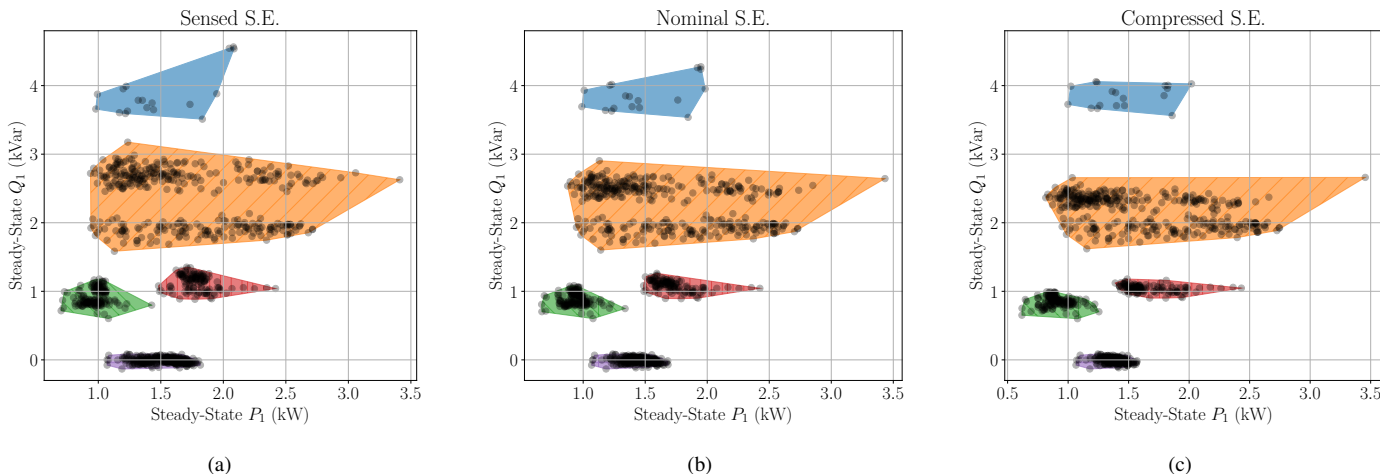


Fig. 11: Electrical events on USCGC Sturgeon across a 15-day period, plotted in the steady-state real and reactive power feature space, using the techniques of Sections IV-A to IV-C. (a) Sensed spectral envelopes. (b) Nominal spectral envelopes. (c) Compressed spectral envelopes.




Fig. 12: Q_1 during the startup transient of the waste vacuum pump on USCGC Sturgeon. Transients are shown for both grid-connected and islanded modes of the ship's microgrid. (a) Sensed spectral envelopes. (b) Compressed spectral envelopes.

0.69, respectively. These scores show that compressed spectral envelopes provide greater feature space separability for this system.

The orange cluster in Fig. 11 corresponds to a waste vacuum pump on USCGC Sturgeon. This pump serves as an example of a load that benefits from using compressed spectral envelopes. In Fig. 11a, this cluster appears to consist of two distinct regions. However in Fig. 11c, preprocessing compresses the two regions into one. The vertical separation into two clusters with sensed spectral envelopes is due to the difference in the ship's voltage amplitude between in-port and underway operation. Specifically, the bottom cluster (at Q_{ss} around 2 kVar) corresponds to the 450 V generator system used while underway, and the top cluster corresponds to the

480 V power source supplied by the terrestrial grid while in port. This pump uses a 3-phase grid-connected induction motor, which has an associated inrush in its startup power transient that serves as a telltale of its operation. Fig. 12 shows the Q_1 stream (computed with sensed and compressed spectral envelopes) for several startup transients of this load. Transients from grid-connected operation are in blue and transients in islanded operation are in orange. In Fig. 12a, the peak and steady-state values of Q_1 are lower in islanded mode than in grid-connected mode, since the islanded voltage amplitude is less than the grid-connected voltage amplitude. However, in Fig. 12b, the two modes have almost identical startup transients due to compressive preprocessing. This example shows the benefit of compressive preprocessing for enabling




Fig. 13: Scatter plot of P and Q of the vacuum pump on USCGC Sturgeon for varying voltage amplitudes. The dashed lines show linear regressions of both P and Q versus the square of the voltage amplitude.




Fig. 14: Vacuum pump electrical events plotted in the steady-state real and reactive power feature space using affine spectral envelopes. Convex hulls are plotted for the corresponding events using this and the other three amplitude compensation techniques.

easier identification of this load from an aggregate power stream.

B. Affine Compensation

Computing affine spectral envelopes for a load using Eqs. (26) and (27) requires an estimate for α_P and α_Q . If these are not known *a priori*, a model built from empirical voltage and power observations from the load can produce them. To illustrate this, Fig. 13 shows a scatter plot of real and

reactive power data (computed with sensed spectral envelopes) versus supply voltage amplitude for the vacuum pump on USCGC Sturgeon. As an induction motor, this load's power consumption contains a constant component. The dashed lines show the results of a linear regression of these P and Q data points as a function of the square of the voltage amplitude (i.e., A^2). These regression lines produce estimates of α_P and α_Q as 3.32 kW and -3.07 kVar, respectively.

Finally, Fig. 14 shows the cluster of vacuum pump electrical events using affine spectral envelopes and the values of α estimated with the linear regression. The red polygon shows the convex hull of this cluster of data. The blue, orange, and green polygons show the corresponding convex hulls to the clusters formed by the other three amplitude scaling techniques. The affine spectral envelope cluster encompasses a much smaller area than the other techniques. Using physical knowledge of the character of a load can therefore produce highly compressed clusters of load data for future pattern recognition and load diagnostics. Similar analysis can be developed for loads with other nonlinear profiles.

VII. CONCLUSION

Modern computing power has made spectral envelope processing and storage fast and cheap. A high-bandwidth power monitor can run a number of preprocessors in real time, each targeted to a different monitoring goal. This "multistream" approach can simplify the task of making power data actionable. Although the case study presented in this work is limited to voltage amplitude and frequency deviations, this work can be extended to many other exogenous variations, such as change in load demand.

ACKNOWLEDGMENTS

We gratefully acknowledge the support and dedication of the US Coast Guard and in particular the crew of USCGC Sturgeon. We also thank Prof. Steven Shaw for the suggestion of calculating impedance with a NILM. We also gratefully acknowledge the support of the Office of Naval Research NEPTUNE program.

REFERENCES

- [1] H. Lu, L. Zhan, Y. Liu, and W. Gao, "A microgrid monitoring system over mobile platforms," *IEEE Transactions on Smart Grid*, vol. 8, no. 2, pp. 749–758, 2017.
- [2] D. Sun, H. Liu, S. Liu, and T. Bi, "Development of synchronized waveform measurement and its application on fault detection," *IEEE Transactions on Instrumentation and Measurement*, vol. 72, pp. 1–11, 2023.
- [3] M. Dong, P. C. M. Meira, W. Xu, and W. Freitas, "An event window based load monitoring technique for smart meters," *IEEE Transactions on Smart Grid*, vol. 3, no. 2, pp. 787–796, 2012.
- [4] B. Cui, A. K. Srivastava, and P. Banerjee, "Synchrophasor-based condition monitoring of instrument transformers using clustering approach," *IEEE Transactions on Smart Grid*, vol. 11, no. 3, pp. 2688–2698, 2020.
- [5] V. Malbasa, C. Zheng, P.-C. Chen, T. Popovic, and M. Kezunovic, "Voltage stability prediction using active machine learning," *IEEE Transactions on Smart Grid*, vol. 8, no. 6, pp. 3117–3124, 2017.
- [6] S. Tang, S. Yuan, and Y. Zhu, "Data preprocessing techniques in convolutional neural network based on fault diagnosis towards rotating machinery," *IEEE Access*, vol. 8, pp. 149 487–149 496, 2020.

- [7] J. Paris, J. S. Donnal, Z. Remscrim, S. B. Leeb, and S. R. Shaw, "The sinefit spectral envelope preprocessor," *IEEE Sensors Journal*, vol. 14, no. 12, pp. 4385–4394, 2014.
- [8] P. A. Schirmer and I. Mporas, "Non-intrusive load monitoring: A review," *IEEE Transactions on Smart Grid*, vol. 14, no. 1, pp. 769–784, 2023.
- [9] S. Welikala, C. Dinesh, M. P. B. Ekanayake, R. I. Godaliyadda, and J. Ekanayake, "Incorporating appliance usage patterns for non-intrusive load monitoring and load forecasting," *IEEE Transactions on Smart Grid*, vol. 10, no. 1, pp. 448–461, 2019.
- [10] H. Çimen, N. Çetinkaya, J. C. Vasquez, and J. M. Guerrero, "A microgrid energy management system based on non-intrusive load monitoring via multitask learning," *IEEE Transactions on Smart Grid*, vol. 12, no. 2, pp. 977–987, 2021.
- [11] A. Aboulhian, D. H. Green, J. F. Switzer, T. J. Kane, G. V. Bredariol, P. Lindahl, J. S. Donnal, and S. B. Leeb, "Nilm dashboard: A power system monitor for electromechanical equipment diagnostics," *IEEE Transactions on Industrial Informatics*, vol. 15, no. 3, pp. 1405–1414, 2019.
- [12] D. H. Green, A. W. Langham, R. A. Agustin, D. W. Quinn, and S. B. Leeb, "Adaptation for automated drift detection in electromechanical machine monitoring," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 34, no. 10, pp. 6768–6782, 2023.
- [13] J. L. Kirtley, "Ac voltage, current, and power," in *Electric Power Principles*. John Wiley & Sons, Ltd, 2020, ch. 2, pp. 17–31.
- [14] J. G. Kassakian, D. J. Perreault, G. C. Verghese, and M. F. Schlecht, *Principles of Power Electronics*, 2nd ed. Cambridge University Press, 2023.
- [15] M. P. Tcheou, L. Lovisolio, M. V. Ribeiro, E. A. B. da Silva, M. A. M. Rodrigues, J. M. T. Romano, and P. S. R. Diniz, "The compression of electric signal waveforms for smart grids: State of the art and future trends," *IEEE Transactions on Smart Grid*, vol. 5, no. 1, pp. 291–302, 2014.
- [16] A. W. Langham, D. H. Green, and S. B. Leeb, "Resolution analysis for power system measurement and transient identification," *IEEE Transactions on Instrumentation and Measurement*, vol. 71, pp. 1–10, 2022.
- [17] S. Welikala, N. Thelasingha, M. Akram, P. B. Ekanayake, R. I. Godaliyadda, and J. B. Ekanayake, "Implementation of a robust real-time non-intrusive load monitoring solution," *Applied Energy*, vol. 238, pp. 1519–1529, 2019.
- [18] P. A. Lindahl, D. H. Green, G. Bredariol, A. Aboulhian, J. S. Donnal, and S. B. Leeb, "Shipboard fault detection through nonintrusive load monitoring: A case study," *IEEE Sensors Journal*, vol. 18, no. 21, pp. 8986–8995, 2018.
- [19] Y. Li, L. Fu, K. Meng, Z. Y. Dong, K. Muttaqi, and W. Du, "Autonomous control strategy for microgrid operating modes smooth transition," *IEEE Access*, vol. 8, pp. 142 159–142 172, 2020.
- [20] F. D'Agostino, A. Fidigatti, E. Ragaini, and F. Silvestro, "Integration of shipboard microgrids within land distribution networks: Employing a ship microgrid to meet critical needs," *IEEE Electrification Magazine*, vol. 7, no. 4, pp. 69–80, 2019.
- [21] N. Doerry, "Naval power systems: Integrated power systems for the continuity of the electrical power supply," *IEEE Electrification Magazine*, vol. 3, no. 2, pp. 12–21, 2015.
- [22] I. C. Patnode, M. J. Bishop, A. W. Langham, D. H. Green, and S. B. Leeb, "Shipboard microgrids and automation," *Naval Engineers Journal*, vol. 135, no. 3, pp. 121–131, 2023.
- [23] M. D. A. Al-Falahi, T. Tarasiuk, S. G. Jayasinghe, Z. Jin, H. Enshaei, and J. M. Guerrero, "Ac ship microgrids: Control and power management optimization," *Energies*, vol. 11, no. 6, 2018.
- [24] P. J. Rousseeuw, "Silhouettes: A graphical aid to the interpretation and validation of cluster analysis," *Journal of Computational and Applied Mathematics*, vol. 20, pp. 53–65, 1987.



Aaron W. Langham received the B.E.E. degree in electrical engineering from Auburn University, Auburn, AL, USA in 2018, and the M.S. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, MA, USA in 2022. He is currently pursuing the Ph.D. degree in electrical engineering and computer science at the Massachusetts Institute of Technology, Cambridge, MA, USA. His research interests include signal processing, machine learning, and computer systems for energy systems.




Thomas C. Krause received the B.S degree in electrical engineering from Purdue University, West Lafayette, IN, USA, in 2019 and the M.S. and E.E. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2021 and 2024, respectively. He is currently pursuing the Ph.D. degree in electrical engineering and computer science with the Massachusetts Institute of Technology.



Daisy H. Green received the B.S. degree in electrical engineering from the University of Hawai'i at Mānoa, Honolulu, HI, USA, in 2015, and the M.S. degree and the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2018 and 2022, respectively. She is currently an assistant professor with the department of electrical and computer engineering at the University of Hawai'i at Mānoa. Her research interests include signal processing, estimation, and machine learning

for energy management and condition monitoring.



Steven B. Leeb received the Ph.D. degree from the Massachusetts Institute of Technology, in 1993. Since 1993, he has been a member on the MIT Faculty with the Department of Electrical Engineering and Computer Science. He also holds a joint appointment with the Department of Mechanical Engineering, MIT. He is concerned with the development of signal processing algorithms for energy and real-time control applications.