Resolution Analysis for Power System Measurement and Transient Identification

Aaron W. Langham, *Graduate Student Member, IEEE*, Daisy H. Green, *Graduate Student Member, IEEE*, and Steven B. Leeb, *Fellow, IEEE*

***Abstract*— Averaging techniques are frequently applied when sampling quasi-periodic signals to enhance signal fidelity. This article investigates the use of spectral envelope preprocessing for evaluating power consumption and harmonic content drawn by ac loads. A statistical characterization of a preprocessor system's resolution given analog input noise is derived. Resolution is defined here as the ability to detect transient step changes in power consumption. With these results, the concept of effective number of bit (ENOB) is extended to this averaging system. Design trade-offs in the number of quantization bits, sampling frequency, noise level, and minimum resolvable transient are explored. The conclusions reached here generalize to any system in which spectral envelopes are extracted.**

***Index Terms*— Digital metering, energy monitoring, power computation, power monitoring, spectral envelopes, transient resolution.**

I. Introduction

A WEALTH of information about the operation of electromechanical devices lies hidden in measurable signals. Many of these signals can be characterized as being locally periodic. That is, at some resolution or level of accuracy, a signal exhibits periodicity, but contains nonperiodic characteristics when observed on a longer time scale. Averaging techniques can increase the accuracy and fidelity of acquisition of these types of signals. Oversampling, both with direct quantization and with noise shaping, can improve resolution at the expense of sample rate, in a form of averaging [1]. Ensemble averaging can make use of analog or "prequantization" noise to reduce quantization error and improve resolution across ensembles of periods [2]. An averaging technique that is particularly useful for the monitoring of electromechanical equipment is spectral envelope preprocessing, which allows for compression of signals sampled at a high sampling rate while retaining salient information about the signal's harmonic content [3]. Spectral envelopes find use in a variety of application domains [4]–[7]. They are particularly useful for power monitoring, where extraction of in-phase and quadrature components is desirable to approximate real and reactive power, respectively. One application of power monitoring is energy disaggregation using a nonintrusive load monitor (NILM), to identify individual load operation from the aggregate meter data. For this application, spectral envelope output data are typically fed into a feature extraction system from which a number of classification systems can be run [8].

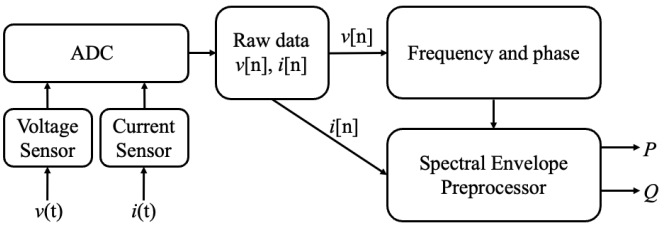

Fig. 1. Conceptual diagram of signal acquisition and preprocessing system.

When calculating spectral envelopes, the voltage and current signals are assumed to be locally periodic over one ac line cycle. At sampling frequency f_s and utility frequency f_0, the average number of samples in one period is $N = f_s/f_0$. Assuming the voltage is well-regulated, or "stiff," the reference voltage signal is $v[n] = V_{\text{pk}} \cdot \sin(2\pi n/N)$, where V_{pk} is the peak voltage magnitude. A "preprocessor" calculates the spectral envelopes by taking local averages of the current signal modulated by a sinusoid over a sliding window, using the voltage signal as the phase and frequency reference [3]. Then, the fundamental real (P) and reactive (Q) power are

$$P = \frac{V_{\text{pk}}}{N} \sum_{n=0}^{N-1} i[n] \cdot \sin(2\pi n/N) \tag{1}$$

$$Q = -\frac{V_{\text{pk}}}{N} \sum_{n=0}^{N-1} i[n] \cdot \cos(2\pi n/N). \tag{2}$$

This is equivalent to the imaginary and real components, respectively, of the discrete Fourier transform (DFT) [3], scaled by $-V_{\text{pk}}/N$. The DFT transforms some signal $i[n]$ into I_k, where k is an integer representing the harmonic order:

$$I_k = \mathcal{F}(i[n]) = \sum_{n=0}^{N-1} i[n] \cdot e^{-jk2\pi n/N}. \tag{3}$$

The analysis presented in this work uses the fundamental ($k = 1$) in-phase component (P) as the preprocessor output without loss of generality due to the orthogonality of sine and cosine. For loads with a non-zero phase angle or with harmonic content, a similar analysis can be done for Q and higher order harmonics that will yield the same insights. An overview of the signal acquisition and preprocessing system is shown in Fig. 1.

Manuscript received October 6, 2021; revised January 13, 2022; accepted February 1, 2022. Date of publication February 11, 2022; date of current version March 2, 2022. This work was supported in part by the Office of Naval Research Naval Enterprise Partnership Teaming with Universities for National Excellence (NEPTUNE) Program and in part by the Grainger Foundation. The Associate Editor coordinating the review process was Dr. Grazia Barchi. *(Corresponding author: Aaron W. Langham.)*

The authors are with the Department of Electrical Engineering and Computer Science (EECS), Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: alangham@mit.edu; dhgreen@mit.edu; sbleeb@mit.edu).

Digital Object Identifier 10.1109/TIM.2022.3151146

Since spectral envelope preprocessing averages or sums over many time points, the number of possible power values is larger than the number of current sample levels that are resolvable with an input quantizer of B bits [3]. However, similar to the quantization artifacts that a quantizer introduces, this averaging is associated with its own set of artifacts and distortions present in the resulting downsampled output signal. The extra preprocessor output values are not evenly spaced across the range of input amplitudes. Thus, the preprocessor's ability to resolve loads (i.e., to both detect and accurately measure step changes) varies based on the input signal. When using a preprocessor output as an input to a feature extraction system for training a classifier, as in the case of a NILM, it is highly undesirable to have training data that contain biases or inaccuracies, even if, or perhaps especially if, they only predominantly affect certain ranges of input signal values.

This work examines the advantages of spectral envelope preprocessing for power measurements and related measurement systems. Design techniques are presented for optimizing the resolution and linearity of spectral envelope data for power monitoring, fault detection, and diagnostic applications. These design techniques are illustrated by contrasting the performance of two different hardware preprocessors. The effects of noise on the power system introduced by data acquisition hardware are included. A new measure of the effective number of bits (ENOBs) or power resolution is developed for the output of a preprocessing system in the presence of prequantization noise, which is illustrated with the data collected on a U.S. Coast Guard (USCG) Cutter. A design guide is developed for the construction or modification of a data acquisition system using a spectral envelope preprocessor, with relevant trade-offs and considerations for parameter choices.

II. Noise-Free Preprocessor Performance

Spectral envelope preprocessing increases the output resolution by averaging over N sampled data points over one or more periods of a quasi-periodic input signal [3]. We begin by considering a a noise-free system with a stiff, sinusoidal voltage. A sinusoidal, in-phase current with amplitude A produces a fundamental frequency spectral envelope associated with real power proportional to A. Quantized sampling results in a discrete number of possible values, 2^B, for each current sample $i[n]$, where B is the number of input quantizer bits. Therefore, over an observation window, there are a finite number of unique sampled waveforms of the current signal, denoted as U [3]. In calculating the spectral envelopes, the DFT, shown in (3), maps each unique input current signal to a corresponding sequence of unique frequency-domain outputs, creating the same number U of unique preprocessor outputs. Effectively, the preprocessor can discern U different waveforms corresponding to different values of A. As derived in [3], the total number of unique preprocessor outputs is

$$U = 2^{B-1} + \sum_{n=1}^{\frac{N}{4}-1} \left\lfloor (2^{B-1} - 1) \cdot \sin\left(\frac{2\pi n}{N}\right) + \frac{1}{2} \right\rfloor. \quad (4)$$

This number of outputs can be characterized in terms of bits by taking $\log_2 U$. For example, for preprocessors with input bits $B = 12$ and $B = 16$ averaged over $N = 133$ data points, the number of output bits is approximately 15 and 19, respectively.

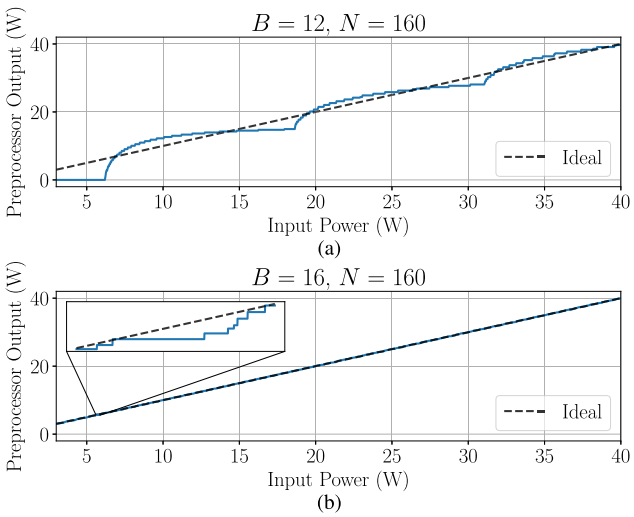

Fig. 2. 12- and 16-bit preprocessor transfer characteristics for a sinusoidal input current. (a) 12-bit preprocessor. (b) 16-bit preprocessor.

Clearly, there are practical resolution benefits from averaging. Current measurements are quantized in a stair-step fashion, with equal spacing between each bit level. The averaging inherent to the preprocessor smooths this quantization. However, the preprocessor "transfer characteristic," the mapping between the actual input power and the preprocessor output level, is nonlinear.

Consider a range of resistive loads powered by a 120-V rms sinusoidal source. The observed current amplitude is swept from approximately 35 to 471 mA, corresponding to a sweep in power consumption from 3 to 40 W. To illustrate the nonlinearity of the spectral envelope calculation, contrast the outputs of two different preprocessors, one operating with 12-bit ($B = 12$) input sampling and the other operating with 16-bit ($B = 16$) sampling, both with $N = 160$ and a maximum quantizable current $I_{\max} = 300$ A (corresponding to approximately 25.5 kW for this system). In this work, all examples use a current sampling rate of 8 kHz; this preprocessor with $N = 160$ thus outputs at 50 Hz. The resulting computations for P from (1) are shown in Fig. 2. The 12-bit preprocessor in Fig. 2(a) shows visible "bumps" and sharp edges forming nonlinearities in the transfer characteristic. The width of the first bump in the plot approximately corresponds to the difference in input power for one least significant bit (LSB) of quantized input current. Fig. 2(b) shows a 16-bit preprocessor that appears more linear at the same scaling. However, the nonlinearities are still present, as shown in the inset of Fig. 2(b). The deviation of each "bump" from an ideal linear transfer characteristic is most noticeable at the lower end of the input power range. The transfer characteristic becomes more linear as the input amplitude is increased. Thus, for measurement of small loads or for preprocessors with relatively few input quantization bits, these nonlinearities or distortions can cause an appreciable bias of the preprocessor output away from the correct value, both for single readings and measured differences in readings.

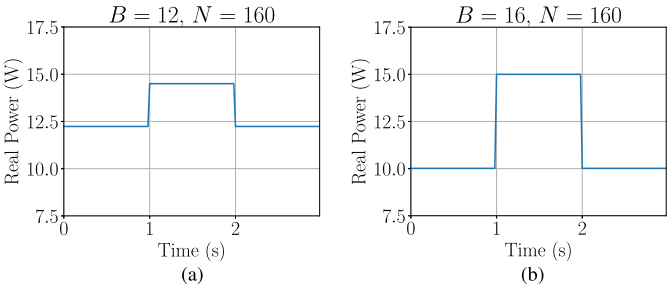

Fig. 3. Simulation of preprocessor nonlinearity for a 5-W load cycling on top of a 10-W base load. (a) 12-bit preprocessor. (b) 16-bit preprocessor.

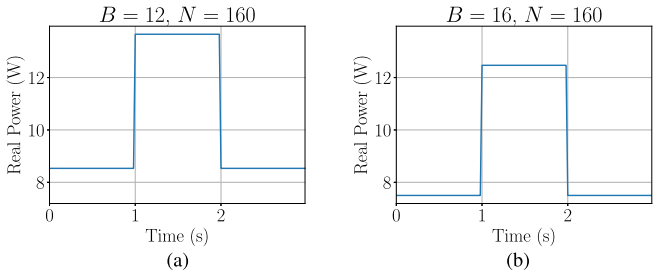

Fig. 4. Preprocessor simulation showing correct differential measurements. (a) 12-bit preprocessor. (b) 16-bit preprocessor.

Deviations from the ideal linear transfer characteristic can introduce significant error in estimating the actual power or, more generally, any spectral envelope, associated with a load. For example, Fig. 3 shows the outputs of the 12- and 16-bit preprocessors for a 5-W resistive load cycling on top of a 10-W base load. The 12-bit preprocessor output is inaccurate at every point in Fig. 3(a). The step change shows a difference of 2.27 W, significantly different than the actual 5-W load demand. That is, even a simple event detector looking at the changes in steady-state power demand would not find the expected 5-W change. The 16-bit preprocessor's output, as shown in Fig. 3(b), is much more accurate. The output value for the base load is 10.01 W, and the output power of the base load and cycling load together is 15 W.

The 12-bit preprocessor is unable to adequately resolve a 5-W difference with the 10-W base load. However, in regions where the slope of the preprocessor output to preprocessor input is approximately linear, the differential change would be closer to the actual 5-W change, even if the measurement at any single point is incorrect. This can be seen in Fig. 4 for a base load of 7.45 W. The values before and after the first step change in Fig. 4(a) are incorrect, at 8.53 and 13.65 W, respectively, but the difference is close to actual, at 5.12 W. For the 16-bit preprocessor, as shown in Fig. 4(b), both the single and differential output measurements are close to the actual values. The values before and after the first step change are 7.5 and 12.47 W, respectively, and the difference is 4.97 W. In a noise-free, idealized environment, the preprocessor's ability to make accurate measurements, both single and differential, depends on the value of the signal being measured.

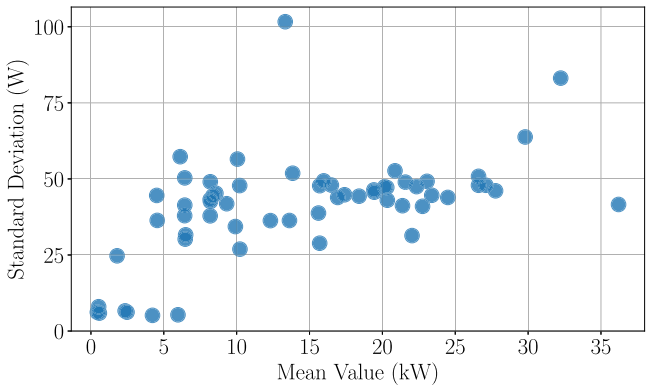

Fig. 5. Preprocessor output standard deviation and mean value for several intervals of shipboard data.

III. Practical Preprocessor Performance

Linearity can be improved by increasing the input bit resolution. However, in a practical system, there are also other factors, such as noise at different stages in the signal processing chain, which can impact the preprocessor performance. Any practical system will contain some non-zero amount of prequantization noise. Due to the uncertainty in each measurement, noise reduces the number of meaningful preprocessor output values. However, prequantization noise can be either a detriment or a benefit to system performance, depending on the amount of noise and preprocessor parameters.

A. Prequantization Noise

Analog distortions in the current signal prior to the quantization stage produce prequantization noise. Prequantization noise arises from many sources, including the thermal noise in the amplifiers of the signal processing chain of the data acquisition (DAQ) hardware [9]. Distortions can also result from electromagnetic coupling in sensors and connections external to DAQ. Another source of distortion is high-frequency current components from the physical operation of other loads on the line, such as nonlinear loads and inductive loads [10], [11]. While not noise in the strict sense, these fluctuations in power drawn by other loads can affect the preprocessor calculations and resolution.

The microgrid of USCG Cutter MARLIN serves as a demonstration. The preprocessor hardware was used to monitor the aggregate current of a subpanel which powers approximately half of the ship. The variance in the preprocessor output for real power (computed as the sum of three phases of real power envelopes) tracks with the mean value of the preprocessor output in steady state. Fig. 5 shows a scatter plot of standard deviation versus mean of real power spectral envelopes for the aggregate shipboard power in steady-state operation. Low-variance values in the bottom left of the scatter plot include resistive loads like heaters. Steady-state windows with higher variance include the operation of loads with significant harmonics or large fluctuations in power draw. An important component of the preprocessor output variance in this industrial environment arises from these fluctuations.

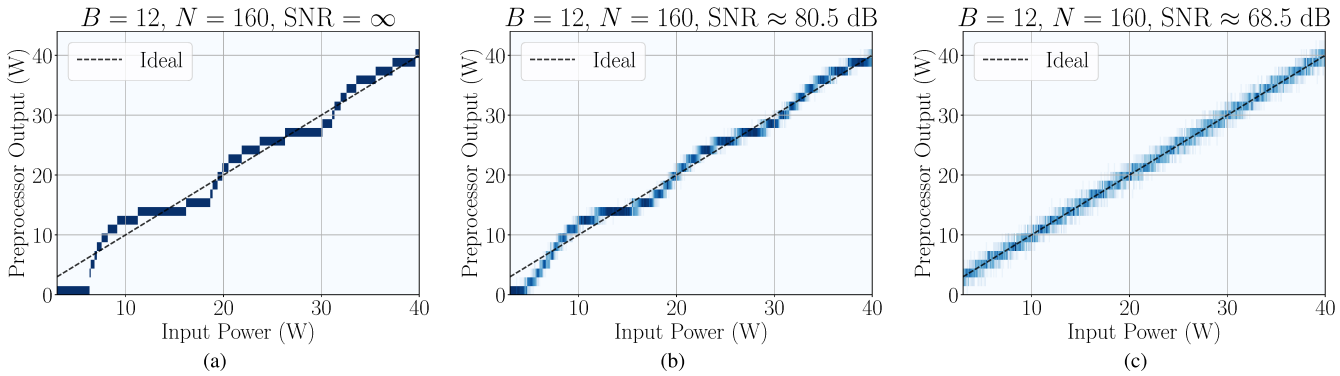

Fig. 6. Preprocessor outputs as histograms for calculated input power values plus Gaussian noise, for a 12-bit quantizer. (a) No prequantization noise. (b) Small amount of prequantization noise. (c) Large amount of prequantization noise.

These sources of prequantization noise are modeled as uncorrelated additive Gaussian white noise (AGWN). The amount of noise introduced into a system is characterized using a signal-to-noise ratio (SNR). Consider a system with a maximum quantizable current of $I_{\max}$ and an rms current noise of σ_G. The maximum quantizable sinusoidal current has rms value $I_{\mathrm{rms}} = I_{\max}/\sqrt{2}$, and the SNR is defined as

$$\mathrm{SNR} = 20\log_{10}\frac{I_{\mathrm{rms}}}{\sigma_G} = 20\log_{10}\frac{I_{\max}}{\sigma_G\sqrt{2}}. \tag{5}$$

B. Transfer Characteristic Linearization

Adding or exploiting a small amount of Gaussian white noise, that is, dither, into an analog-to-digital converter or quantization system linearizes the mean output value of the quantizer [12]. The beneficial dithering effect can improve the accuracy of an acquired current signal and can decrease error in the preprocessor output [3]. The dithering benefit of prequantization noise "propagates" through the preprocessor system, and as a result, the relationship of the mean output versus input gradually becomes more linear as noise is added to the system. This effect is shown in Fig. 6, where prequantization noise is gradually added to a system with $B = 12$ and $N = 160$. Here, the preprocessor transfer characteristic is shown as a series of vertical histograms for a range of input values, with the average of each histogram forming an "average transfer characteristic." The highly nonlinear average transfer characteristic of Fig. 6(a) becomes more linear when a small amount of prequantization noise is added in Fig. 6(b) and becomes even more linear when sufficient noise is added in Fig. 6(c). The preprocessor average transfer characteristic becomes linear at an rms noise value of approximately one half of the LSB of the input quantizer. In Fig. 7, a vertical cross section of Fig. 6(c) is shown for an input power of 18.6 W, corresponding to a point where the transfer characteristic is highly nonlinear in the noise-free preprocessor of Fig. 2(a). A Gaussian distribution is fit with the calculated mean and standard deviation. As can be seen in the plot, the mean output from this preprocessor is approximately equal to the input power of 18.6 W. However, the correct mean value does not come without drawbacks. The variance from noise increases the uncertainty in measurement, as evident by the higher vertical spread of the histograms of Fig. 6(b) and (c). The measurements of the preprocessor output will therefore appear to be "noisy."

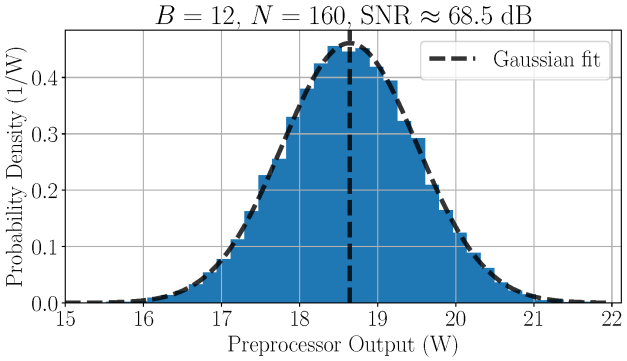

Fig. 7. Histogram of preprocessor outputs for an input power of 18.6 W with the same parameters as Fig. 6(c).

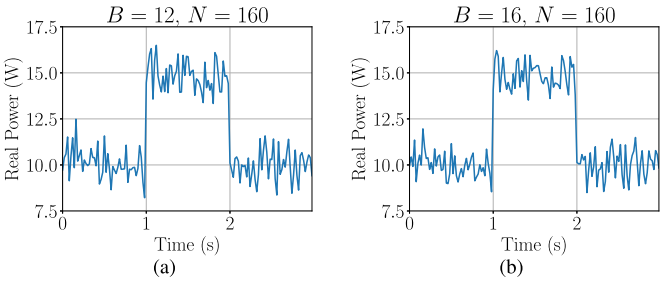

Fig. 8. Simulation of preprocessor linearity from noise for a small load while a base load is running. (a) 12-bit preprocessor. (b) 16-bit preprocessor.

The same load cycle as in Fig. 3 is simulated again, but this time 80 mA rms of Gaussian current noise is added, for an SNR of 68.5 dB [the same amount of noise as in Fig. 6(c)]. The resulting P streams are shown in Fig. 8. For the 12-bit preprocessor in Fig. 8(a), the average value before the step change is around 10.14 W and the average value after the first step change, between 1 and 2 s, is around 14.76 W. The differential measurement is much closer to 5 W, the true power consumption of the load, than the result in Fig. 3(a). The noise linearizes the transfer characteristic and leads to more accurate average single and differential measurements. For the 16-bit preprocessor in Fig. 8(b), these

values are around 10.06 and 15.24 W, respectively. Here, the average single and differential measurements are similar to those observed in Fig. 3(b). The noise-free 16-bit preprocessor transfer characteristic was already close to linear. Thus, adding noise did not significantly change the average, but it did create more uncertainty in any single output power measurement.

C. Effective Number of Bits

A reasonable level of input prequantization noise helps linearize the preprocessor performance. More than this level degrades the output estimates with undesirable noise. To quantify the effect of prequantization noise on resolution ability, the system's "effective number of bits" is examined. Per IEEE standard 1057 [13], effective number of bits (ENOB) is defined as the number of bits of a theoretical quantizer whose rms quantization noise equals the rms noise of the system in question. For some system yielding output x, the ENOB is

$$\text{ENOB} = \log_2\left(\frac{x_{\text{pk-pk}}}{\sigma_x\sqrt{12}}\right) \tag{6}$$

where $x_{\text{pk-pk}}$ is the full range of the system output, and σ_x is the rms noise and distortion of x. In this work, a liberal use of this definition will allow a derivation of the ENOB of the spectral envelope preprocessor output, given some amount of prequantization noise. With the results of this derivation, the effect of noise on the performance of the preprocessor in terms of its number of meaningful "bits" can be examined.

The goal is to find both $P_{\text{pk-pk}}$, the full-scale swing possible in preprocessor output, and σ_P, the standard deviation in the preprocessor output due to input noise. For this derivation, σ_P is assumed to be constant across all input power levels. The input current signal is now treated as a random process, $I[n]$. Using (1), the equation for the fundamental real power spectral envelope (P), for a window size of N, can be written as

$$P = \frac{V_{\text{pk}}}{N}\sum_{n=0}^{N-1} I[n]\cdot\sin(2\pi n/N). \tag{7}$$

It is further assumed that the input noise, represented here as $I_g[n]$, is white and Gaussian and has zero mean and some variance σ_G^2. To show the effect of current signal quantization on the preprocessor input, another assumption is made that the distortion from quantization, represented as $I_q[n]$, can be approximated as white uniform noise with zero mean and variance σ_Q^2 [1]. For a quantizer that operates between $-I_{\max}$ and $+I_{\max}$

$$\sigma_Q^2 = \frac{\Delta^2}{12}, \quad \Delta = \frac{I_{\max}}{2^{B-1}} \tag{8}$$

where Δ is the size of the LSB. These approximations are suitable when the average preprocessor transfer characteristic is well-approximated as linear, such as in Fig. 6(c). For convenience in derivation, the sum of $I_g[n]$ and $I_q[n]$ will be represented as $I_o[n]$, which will also be white noise with zero mean and variance $\sigma_o^2 = \sigma_G^2 + \sigma_Q^2$ (since $I_g[n]$ and $I_q[n]$ are assumed to be independent). Thus, the full representation of $I[n]$ is as follows:

$$I[n] = i_o[n] + I_g[n] + I_q[n] = i_o[n] + I_o[n] \tag{9}$$

where $i_o[n]$ is the deterministic input current signal. Since P is now a random variable, its variance, σ_P^2, can be computed as $E\{P^2\} - E\{P\}^2$, where $E\{\cdot\}$ represents the expected value of a random variable [14]. Due to the linearity of the expected value operator and the fact that $I_o[n]$ is zero-mean, $E\{P\}^2$ can be computed as

$$E\{P\}^2 = \frac{V_{\text{pk}}^2}{N^2}\left(\sum_{n=0}^{N-1}\sin\left(\frac{2\pi n}{N}\right) i_o[n]\right)^2. \tag{10}$$

Next, $E\{P^2\}$ can be found by taking the expected value of the expansion of P^2 and recognizing that the expected value of the random cross term will be zero, since it is zero-mean. The expected value of the deterministic square term will be the value of $E\{P^2\}$ found in (10). Thus, σ_P^2 will simply be the expected value of the random square term. By rewriting it as a product of sums with different indices and combining the sums, the following expression for σ_P^2 can be found:

$$\sigma_P^2 = \frac{V_{\text{pk}}^2}{N^2}\sum_{m=0}^{N-1}\sum_{n=0}^{N-1}\sin\left(\frac{2\pi m}{N}\right)\sin\left(\frac{2\pi n}{N}\right) E\{I_o[m]I_o[n]\}. \tag{11}$$

Because $I_o[n]$ was assumed to be white noise, the autocorrelation of $I_o[n]$ is zero except for the zero-shift case, and $E\{I_o[m]I_o[n]\}$ becomes $\sigma_o^2\delta[m-n]$. Through the filtering property of the delta function, the m summation and the delta function can be discarded and m can be set equal to n:

$$\sigma_P^2 = \frac{V_{\text{pk}}^2\cdot\sigma_o^2}{N^2}\sum_{n=0}^{N-1}\sin^2(2\pi n/N). \tag{12}$$

Using the power reduction identity and Lagrange's trigonometric identity [15], it can be shown that for any integer $N > 2$, the sum above will simplify to $N/2$. Thus, the expression for σ_P becomes the following, since $V_{\text{rms}} = V_{\text{pk}}/\sqrt{2}$:

$$\sigma_P = \frac{V_{\text{rms}}\cdot\sigma_o}{\sqrt{N}} = \frac{V_{\text{rms}}\sqrt{\sigma_G^2+\sigma_Q^2}}{\sqrt{N}}, \quad N > 2. \tag{13}$$

This implies that for some input current noise, the corresponding preprocessor output noise will be the product of rms voltage and rms current noise, but reduced as the number of averaging points increases. This makes intuitive sense, as in many statistical contexts variance tends to decrease in proportion to the number of samples (N) and the standard deviation tends to decrease proportionally with $\sqrt{N}$. Next, it is assumed that the full-scale preprocessor swing, $P_{\text{pk-pk}}$, is approximately equal to $V_{\text{pk}}\cdot I_{\max}$, i.e., twice the maximum real power observable, occurring when there is zero phase angle between the voltage and current.

With σ_P and $P_{\text{pk-pk}}$ obtained, these quantities can be used in the equation for ENOBs given in (6). To characterize noise as the third parameter in the preprocessing system (alongside B and N), a new parameter G is defined as the linear form of the quantizer's SNR, as defined in (5)

$$G = \left(\frac{I_{\text{rms}}}{\sigma_G}\right)^2 = \frac{1}{2}\left(\frac{I_{\max}}{\sigma_G}\right)^2 = 10^{\text{SNR}/10} \tag{14}$$

where G is then substituted into (13). σ_Q can be substituted out of this equation and replaced with $\Delta/\sqrt{12}$ due to the assumption made previously about the character of the quantization noise. Since $\Delta = I_{\max}/2^{B-1}$ from (8), it can be substituted out of the equation for ENOB. The product and quotient rules for logarithms can also be used, and the terms can be canceled and rearranged to obtain the final representation of ENOB of the preprocessor for a given B, N, and G

$$\text{ENOB} = B + \frac{1}{2}\log_2\frac{N}{2} - \frac{1}{2}\log_2\left(\frac{3\cdot 4^B}{2G} + 1\right), \quad N > 2. \tag{15}$$

This equation shows that for the ENOB of the preprocessor, the number of bits from the quantizer (B) is the baseline, and then some benefit is obtained from increasing the number of averaging points (N), but only logarithmically. If there is no noise, that is, if $G = \infty$ (infinite SNR), then the third term goes to zero, meaning that there is no penalty in the ENOB from prequantization noise. However, if there is prequantization noise, it will lower the ENOB of the preprocessor approximately logarithmically.

IV. Demonstration With Hardware

Spectral envelope preprocessors using two different data acquisition boards are examined in this section. The first is the LabJack UE9 data acquisition hardware, which obtains data at a 12-bit input resolution. The second is the data acquisition hardware of [16], which obtains data at a 16-bit input resolution. Both boards sample voltage and current signals at a 8-kHz sampling rate, for each phase of the power system. One period is used for the spectral envelope window with no window overlap, so the output stream has a sampling rate equal to the line frequency (60 Hz). Thus, the average value of N is $133.\bar{3}$, or approximately 133. Both DAQs' signal processing chains introduce approximately half an LSB rms of prequantization noise into their preprocessors, resulting in linear average transfer characteristics. Due to the difference in the number of bits, this means that the 16-bit DAQ contains less noise than the 12-bit DAQ. The current sensor hardware is the LEM LF-305 which has a maximum current limit of 300 A. The quantizer is chosen such that its maximum quantizable value, referred to here as $I_{\max}$, is aligned with the maximum value of the sensor (i.e., $I_{\max} = 300$ A). If $I_{\max}$ is larger than the maximum sensor value, there will be a degradation in resolution since the LSB will be unnecessarily large; if $I_{\max}$ is smaller than the maximum sensor value, there will be clipping in the input signal if a value over $I_{\max}$ is acquired.

In Fig. 9, the ENOB for the preprocessor output versus SNR, as derived in (15), has been plotted for the case of 12- and 16-bit preprocessors with $N = 133$. Three operating points are marked. The leftmost operating point is for both the 12- and 16-bit DAQs monitoring the starboard-side aggregate shipboard power of USCG Cutter MARLIN (in this context, aggregate refers to metering the entire panel and its collection of loads, as opposed to submetering a smaller branch of the power delivery system). For this point, the ENOB of the preprocessor output was obtained by approximating the standard deviation of the steady-state preprocessor output. Since there is so much noise, the calculated ENOB of both these DAQs is approximately 10.4, meaning the 16-bit DAQ holds no significant advantage here. A comparison of the two DAQs' power streams can be seen in Fig. 10 for the starboard-side subpanel of USCG Cutter MARLIN. In Fig. 10(a) and (b), for the 12- and 16-bit DAQs, respectively, an approximately 200-W resistive load is energized on top of the aggregate shipboard power (at different times, hence at different base loads), turning on at around $t = 0.5$ s and cycling three times for 1-s ON and 2-s OFF. The load is equally resolvable for both the 12- and 16-bit DAQs. In Fig. 10(c) and (d), an approximately 15-W resistive load is energized on top of the aggregate shipboard power, for the 12- and 16-bit DAQs, respectively. The 15-W resistive load cycles three times, as in the previous scenario; however it is not visible for either DAQ, because of the large fluctuations in power drawn by other loads.

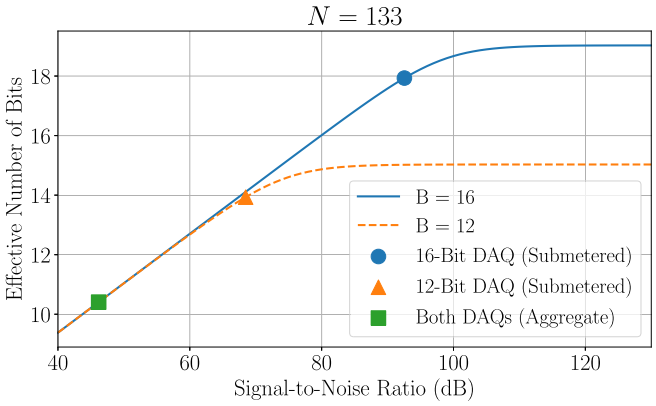

Fig. 9. Simulated ENOB for 16- and 12-bit preprocessors for a sweep of SNR values. Also plotted are the ENOB from measurements of two DAQs, each in an aggregate (i.e., metering the entire subpanel) and submetered (i.e., metering a smaller subset of the panel) environment. In the noisy environment (the entire subpanel), both DAQs have the same ENOB (green square).

In Fig. 9, moving rightward from the leftmost operating point, as the amount of noise decreases (SNR increases), the ENOB of each preprocessor increases essentially linearly until it reaches a point of saturation. For the 12-bit DAQ, the ENOB saturates at around 15 bits, and for the 16-bit DAQ, the ENOB saturates at around 19 bits. Note that these saturated values approximately match the values calculated in Section II using the formula provided in [3].

A second shipboard setup was configured to monitor an individually submetered outlet with the same 12- and 16-bit DAQs used previously. In this case, the dominant prequantization noise is the thermal noise in the analog processing chain, since no other load currents are being metered. The middle operating point in Fig. 9 is for the 12-bit DAQ in a submetered environment with an internal noise of $\sigma_G \approx 80$ mA (SNR ≈ 68.5 dB). The rightmost point is for the 16-bit DAQ in a submetered environment, with internal prequantization noise of $\sigma_G \approx 5$ mA (SNR ≈ 92.6 dB). The prequantization noise values can be used to obtain ENOB using (15). The middle and rightmost operating points yield an ENOB of approximately

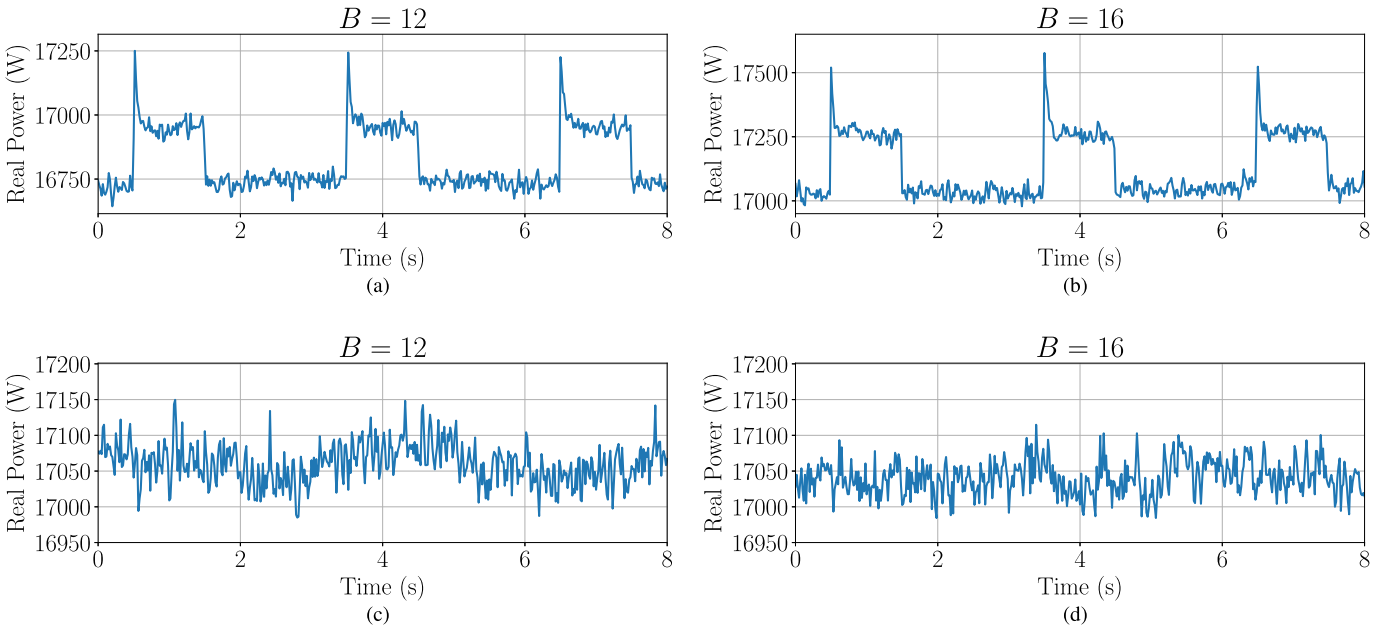

Fig. 10. 12- and 16-bit preprocessor outputs in an aggregate shipboard microgrid for 200- and 15-W loads. (a) 200-W load on the aggregate power stream (12-bit DAQ). (b) 200-W load on the aggregate power stream (16-bit DAQ). (c) 15-W load on the aggregate power stream (12-bit DAQ). (d) 15-W load on the aggregate power stream (16-bit DAQ).

14 and 18, respectively. These are only slightly lower than the maximum saturated value, due to the low amount of noise.

V. Design Guide

The design of a spectral envelope preprocessing system requires consideration of three parameters: number of quantization bits, B; window length, N; and signal-to-noise ratio, SNR. The minimum resolvable load of the target application is another parameter that is either tunable or fixed for a given system. The minimum resolvable load can be expressed either with ENOB or with a ratio of minimum to maximum resolvable powers. This section guides the design of a spectral envelope preprocessing system based on the parameters that are either tunable or fixed for a particular system.

A. Measuring System Noise

The preprocessor quantizer's SNR is the ratio of the square of the rms maximum quantizable input current (I_{rms}^2) to the variance of the prequantization noise in the current signal (σ_G^2), as defined in (5). To determine the SNR, the distribution of noise is approximated as a zero-mean Gaussian, whose variance must be either acquired or designed. If the variance is not known *a priori*, it can be approximated in multiple ways. If the current signal can be measured with no loads drawing power, the standard deviation of this signal can be computed over a window. If the current signal data are only available for an energized system, a window of the current signal for steady-state power consumption can be obtained, and then a three-parameter sine-wave fitting algorithm can be applied [17]. The generated sinusoid can then be subtracted from the current signal and the standard deviation obtained from this resulting difference signal. This method has the advantage that if the monitored system contains variance that is not independent of the loads energized, such as that in Fig. 5, it can find an approximation of the standard deviation for any load configuration. Finally, if the preprocessor transfer characteristic can be assumed to be linear due to sufficient prequantization noise, the standard deviation of the preprocessor output (σ_P) over a window can be calculated. Then, the standard deviation of the quantized current signal (σ_G) can be approximated by rearranging (13)

$$\sigma_G = \sqrt{\frac{\sigma_P^2 N}{V_{\text{rms}}^2} - \sigma_Q^2}. \tag{16}$$

B. Choosing the Number of Bits

As shown previously in Fig. 10, increasing the number of quantizer bits in the presence of excessive noise may be fruitless. Thus, when designing a system to operate in noisy conditions, the knowledge of SNR and window length N can be used with (15) to compare various candidate values for B and assess whether the marginal change in ENOBs of the preprocessor is worth the extra expense of the increase in B. Due to the approximations made in the derivation of this equation, it should only be used when the preprocessor average transfer characteristic can be assumed to be approximately linear due to sufficient prequantization noise. For the case when the preprocessor transfer characteristic is nonlinear, increasing B will result in an appreciable benefit in the preprocessor's ability to resolve small loads by reducing the size of the nonlinear regions.

For example, Fig. 11 shows a comparison of four different preprocessors for a nominal 7.5-W load energized on top of an approximately 2.5-W base load (corresponding to the smart plug used for cycling the load) in the submetered environment on USCG Cutter MARLIN. Fig. 11(a) shows the resulting power stream for the 12-bit DAQ. The prequantization noise linearizes the average preprocessor output,

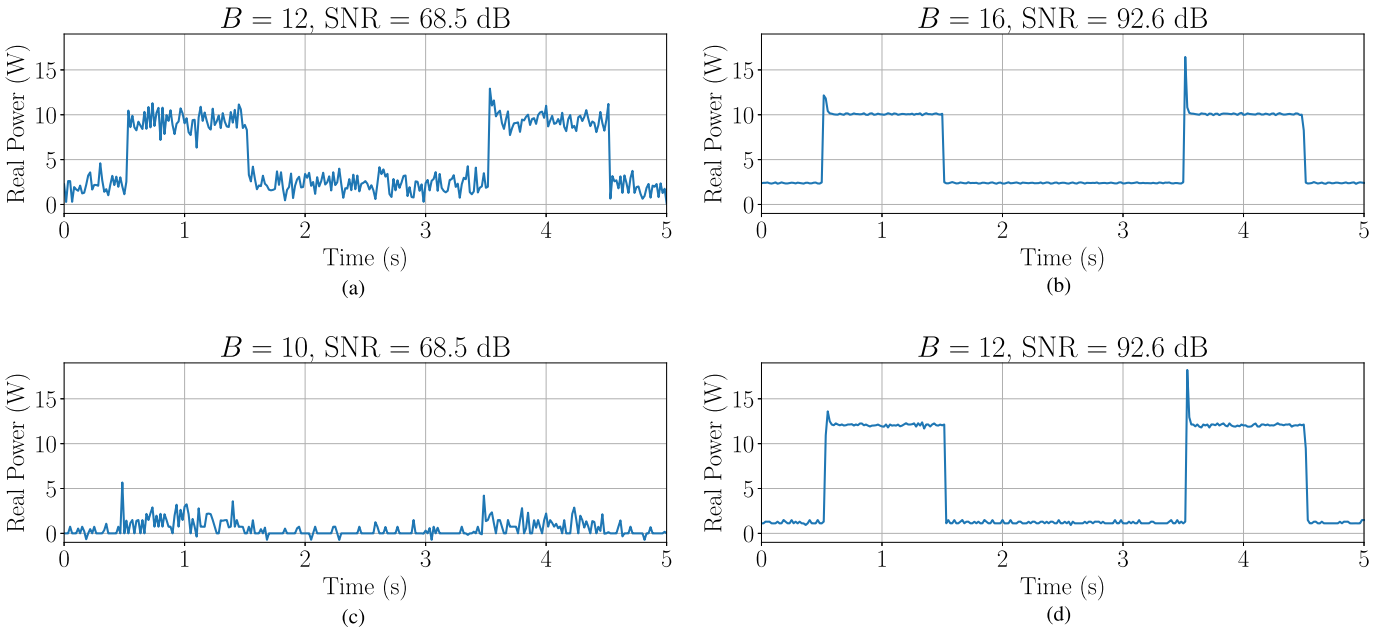

Fig. 11. Comparison of four different preprocessors for a nominal 7.5-W load turn-on on an approximately 2.5-W base load. (a) 12-bit DAQ. (b) 16-bit DAQ. (c) 12-bit DAQ truncated to 10 bits. (d) 16-bit DAQ truncated to 12 bits.

leading to a difference in the average steady-state values of 7.06 W, close to the correct value. Using the 16-bit DAQ, as shown in Fig. 11(b), the steady-state and difference in steady-state power values are correct. The difference in steady-state when the load turns on is approximately 7.5 W. Then, the number of bits is artificially reduced through truncation for both DAQs. Fig. 11(c) shows the power stream for the 12-bit DAQ downquantized to $B = 10$. In this case, the load becomes barely visible. Reducing the 16-bit DAQ to $B = 12$ results in a deterioration of accuracy, as shown in Fig. 11(d). Due to low noise, the transfer characteristic of this case is nonlinear, and the difference in power values when the load is turned on is approximately 10.8 W, which is clearly incorrect.

C. Choosing the Sampling Rate

The parameter N governs the number of data points in the spectral envelope averaging window. For the case of a window length of one period, N is the ratio of sampling frequency to operating frequency. Broadly, N controls the shape and smoothness of the transfer characteristic nonlinearities compared with the "stair-step" transfer characteristic of a simple quantizer. However, the marginal benefit of increasing N drops off logarithmically. If the sampling frequency is not an integer multiple of the operating frequency, then N will alternate between different values for each spectral envelope calculation window. For example, for a sampling frequency of 8 kHz and an operating frequency of 60 Hz, every third window will have an N of 134, and two out of every three windows will have an N of 133, resulting in an average N of $133.\overline{3}$. As a result of inconsistent window sizes, the preprocessor output for this situation will appear to have "spikes" every third sample, even for what should be a constant value. This can be removed using a rolling average of length three, at the expense of smoothing the transients and potentially reducing the ability to identify load transients. These fluctuations may be inconsequential if there is a high amount of variance in the signal. In this work, this effect is avoided in the simulations using a 50-Hz utility frequency, and it is ignored for the 60-Hz demonstrations since the fluctuations are negligible.

D. Ability to Resolve Small Loads

For a given system installation, there exists some minimum load power consumption that can be resolved. Spectral envelope preprocessing introduces two issues that impact the ability of a system to resolve small loads. First, the nonlinearity of the noise-free preprocessor transfer characteristic results in a reduction in resolution at certain input values, as shown in Fig. 3. Second, significant prequantization noise reduces the number of meaningful output values even with an increase in the number of quantizer bits. For both the cases, it is useful to examine the preprocessor's ability to differentiate a small load event given values for B, N, and some base load.

If the prequantization noise is sufficient to linearize the average preprocessor transfer characteristic, the minimum resolvable power will be independent of any base load value. The minimum power that can be resolved by the preprocessor, $P_{\min}$, can be estimated as σ_P. The linear signal-to-noise ratio G from (14) can be related to σ_P from (13). $I_{\max}/\sqrt{2G}$ can be substituted for σ_G and $I_{\max}/(2^{B-1}\sqrt{12})$ for σ_Q. Since the maximum power detectable by the system can be approximated as $P_{\max} = V_{\text{rms}} \cdot I_{\text{rms}}$, where $I_{\text{rms}} = I_{\max}/\sqrt{2}$, the following equation for $P_{\min}/P_{\max}$ can be derived:

$$\frac{P_{\min}}{P_{\max}} = \sqrt{\frac{1}{N}\left(\frac{1}{G} + \frac{2}{3 \cdot 4^B}\right)}. \tag{17}$$

For example, this can be calculated for the USCG Cutter MARLIN example in Fig. 10(b) over a window between $t = 1.75$ s and $t = 3.25$ s. The standard deviation of the real power preprocessor output is computed to be approximately 22.7 W, which will be used as $P_{\min}$. Since the rms voltage is 254 V and the maximum quantizable current signal is 300 A, $P_{\max}$ is found to be 53881.5 W. Thus, the ratio of $P_{\min}$ to $P_{\max}$ is $1:2374$. This implies that all other parameters being constant, increasing the maximum observable power (likely by changing $I_{\max}$) will increase the minimum observable load power by the same factor, providing a heuristic for the range of observable values. $P_{\min}/P_{\max}$ can be related to ENOB through the following equation:

$$\text{ENOB} = -\log_2\left(\sqrt{3}\cdot\frac{P_{\min}}{P_{\max}}\right). \tag{18}$$

Either $P_{\min}/P_{\max}$ or ENOB can be used to quantify the ability of a system to resolve small loads.

If prequantization noise does not linearize the average preprocessor transfer characteristic, the system's resolution will be lower at the nonlinear regions of the transfer characteristic. For example, in Fig. 2(a), between an input power of approximately 12–18 W, there will be very little difference in the output value assigned. What counts as "resolvable" is application-specific. For example, consider a correlation matching algorithm with a set tolerance [18]. If a match tolerance of 20% is chosen, a load event is considered to be resolvable if the difference in preprocessor outputs is within 20% of the actual load value. Considering the base and small load scenario, different base load values can be iterated, and for each, the minimum small load value is found that yields a preprocessor output difference within the desired tolerance. Finding a closed-form solution for this value for any given match tolerance is not tractable, but in general, the worst case is at the first inflection point, when the preprocessor output becomes non-zero [e.g., at approximately 6.2 W in the Fig. 2(a)]. For power monitoring applications, it is unlikely that the preprocessor would operate in this nonlinear situation, as only approximately half of an LSB of rms noise current signal is required to linearize the average transfer characteristic.

E. Design Scenarios

The presented techniques can inform system design. Given three fixed or desired parameters, (17) can be used to solve for the variable of interest. In the first scenario, N and SNR are fixed and there is some requirement on the minimum resolvable load, which fixes $P_{\min}/P_{\max}$. This scenario could be brought about by design of data acquisition hardware at some required sampling frequency in an environment with a known amount of noise. Here, solving for B will yield the number of input quantization bits necessary to meet the resolution requirement. However, it is not always possible to meet the desired requirements by only tuning B, as there is no realizable value of B that will yield a $P_{\min}/P_{\max}$ of less than $(1/(NG))^{1/2}$. In the second scenario, N and B are fixed, perhaps due to use of specified data acquisition hardware, and $P_{\min}/P_{\max}$ is set to some desired value. Solving for the minimum SNR for such a resolution requirement allows a system designer to evaluate potential sensor placements. That is, can a sensor be placed "upstream" to monitor a larger and most likely noisier system, or are multiple "downstream" sensors required? Similar to the first scenario, $P_{\min}/P_{\max}$ is lower bounded. In this scenario, the lower bound is $(2/(3\cdot 4^B\cdot N))^{1/2}$, so even an infinite SNR will not meet every requirement. Finally, in the third scenario, N and B are again fixed due to prespecified hardware. SNR is also fixed in this scenario, which may be due to unavoidable noise in the installation environment. By solving for $P_{\min}/P_{\max}$, the designer can understand which loads can and cannot be resolved in this installation.

VI. Conclusion

Spectral envelope preprocessing can enhance system resolution beyond that of the input quantizer. Prequantization noise can be either a benefit or a drawback depending on the amount introduced. A lack of prequantization noise will result in a nonlinear preprocessor transfer characteristic. An excessive amount of prequantization noise will cause preprocessor systems with different numbers of input quantization bits to operate at similarly poor effective resolutions. By adapting the concept of effective number of bits to the spectral envelope preprocessing system, this effect can be evaluated quantitatively for a proposed preprocessor design given the number of quantizer bits, number of points in the averaging window, and signal-to-noise ratio. The ratio of minimum resolvable power to maximum resolvable power can be derived for the preprocessor configuration to describe the system's ability to resolve small loads.

Acknowledgment

The authors would like to thank the U.S. Coast Guard and in particular the crew of USCGC MARLIN for granting access to their ship.

References

[1] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing*, 3rd ed. Upper Saddle River, NJ: Pearson, 2010.

[2] R. Skartlien and L. Oyehaug, "Quantization error and resolution in ensemble averaged data with noise," *IEEE Trans. Instrum. Meas.*, vol. 54, no. 3, pp. 1303–1312, Jun. 2005.

[3] J. Paris, J. S. Donnal, Z. Remscrim, S. B. Leeb, and S. R. Shaw, "The sinefit spectral envelope preprocessor," *IEEE Sensors J.*, vol. 14, no. 12, pp. 4385–4394, Dec. 2014.

[4] J. A. de la O Serna, W. Van Moer, and K. Barbé, "Using alternating Kalman filtering to analyze oscillometric blood pressure waveforms," *IEEE Trans. Instrum. Meas.*, vol. 62, no. 10, pp. 2621–2628, Oct. 2013.

[5] H. Kameoka, N. Ono, and S. Sagayama, "Speech spectrum modeling for joint estimation of spectral envelope and fundamental frequency," *IEEE Trans. Audio, Speech, Language Process.*, vol. 18, no. 6, pp. 1507–1516, Aug. 2010.

[6] S. R. Shaw and C. R. Laughman, "A Kalman-filter spectral envelope preprocessor," *IEEE Trans. Instrum. Meas.*, vol. 56, no. 5, pp. 2010–2017, Oct. 2007.

[7] D. S. Stoffer, D. E. Tyler, and D. A. Wendt, "The spectral envelope and its applications," *Stat. Sci.*, vol. 15, no. 3, pp. 224–253, Aug. 2000.

[8] D. H. Green, S. R. Shaw, P. Lindahl, T. J. Kane, J. S. Donnal, and S. B. Leeb, "A MultiScale framework for nonintrusive load identification," *IEEE Trans. Ind. Informat.*, vol. 16, no. 2, pp. 992–1002, Feb. 2020.

[9] *Fundamentals of Precision ADC Noise Analysis*, Texas Instruments, Dallas, TX, USA, Sep. 2020. [Online]. Available: https://www.ti.com/lit/pdf/SLYY192

[10] D. Kumar and F. Zare, "A comprehensive review of maritime microgrids: System architectures, energy efficiency, power quality, and regulations," *IEEE Access*, vol. 7, pp. 67249–67277, 2019.
[11] A. Testa *et al.*, "Interharmonics: Theory and modeling," *IEEE Trans. Power Del.*, vol. 22, no. 4, pp. 2335–2348, Oct. 2007.
[12] L. Schuchman, "Dither signals and their effect on quantization noise," *IEEE Trans. Commun. Technol.*, vol. COM-12, no. 4, pp. 162–165, Dec. 1964.
[13] *IEEE Standard for Digitizing Waveform Recorders*, Standard 1057-2017 (Revision of IEEE Std 1057-2007), 2018, p. 1.
[14] M. Loève, "Probability theory I," *Graduate Texts Math.*, vol. 45, p. 12, Jul. 1977.
[15] E. O. Muñiz, "A method for deriving various formulas in electrostatics and electromagnetism using Lagrange's trigonometric identities," *Amer. J. Phys.*, vol. 21, no. 2, pp. 140–141, Feb. 1953.
[16] K. Lee, L. Huchel, D. H. Green, and S. B. Leeb, "Automatic power frequency rejection instrumentation for nonintrusive frequency signature tracking," *IEEE Trans. Instrum. Meas.*, vol. 70, pp. 1–11, 2021.
[17] *IEEE Standard for Terminology and Test Methods for Analog-to-Digital Converters*, Standard 1241-2010 (Revision of IEEE Std 1241-2000), 2011, pp. 1–139.
[18] J. Paris, J. S. Donnal, and S. B. Leeb, "NilmDB: The non-intrusive load monitor database," *IEEE Trans. Smart Grid*, vol. 5, no. 5, pp. 2459–2467, Sep. 2014.

Aaron W. Langham (Graduate Student Member, IEEE) received the B.E.E. degree in electrical engineering from Auburn University, Auburn, AL, USA, in 2018. He is currently pursuing the Ph.D. degree in electrical engineering and computer science with the Massachusetts Institute of Technology, Cambridge, MA, USA.

His research interests include signal processing, machine learning, and computer systems for energy management.

Daisy H. Green (Graduate Student Member, IEEE) received the B.S. degree in electrical engineering from the University of Hawaii at Mānoa, Honolulu, HI, USA, in 2015, and the M.S. degree in electrical engineering from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2018, where she is currently pursuing the Ph.D. degree.

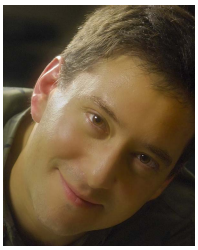

Steven B. Leeb (Fellow, IEEE) received the Ph.D. degree from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 1993.

Since 1993, he has been a member on the MIT Faculty with the Department of Electrical Engineering and Computer Science. He currently 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.