

Transient Event Detection in Spectral Envelope Estimates for Nonintrusive Load Monitoring

Steven B. Leeb, Member, IEEE Steven R. Shaw James L. Kirtley, Jr., Fellow, IEEE

Laboratory for Electromagnetic and Electronic Systems
Massachusetts Institute of Technology
Cambridge, MA 02139, USA

Abstract: *This paper describes the theoretical foundation and prototype implementation of a transient event detector for use in a nonintrusive load monitor (NILM). The NILM determines the operating schedule of the major electrical loads in a building from measurements made at the electric utility service entry. The transient event detector extends the applicability of the NILM to challenging commercial and industrial sites. A spectral preprocessor for use in the transient event detector is introduced first. Then, the transient event detection algorithm is developed. The performance of the algorithm is illustrated with results from a prototype event detector.*

I. Background

The transient behavior of a typical electrical load is strongly influenced by the physical task that the load performs. The load survey conducted in [1] indicates that intrinsic properties modeled as nonlinearities in the constitutive laws of the elements that comprise a load, or in the state equations that describe a load, or both, create repeatably observable turn-on transient profiles suitable for identifying specific load classes. The turn-on transients associated with a fluorescent lamp and an induction motor, for example, are distinct because the physical tasks of igniting an illuminating arc and accelerating a rotor are fundamentally different. Distinctive transient profiles tend to persist even in loads which employ active waveshaping or power factor correction.

This observation has led to the development of a transient event detector for nonintrusive load monitoring.

95 WM 042-2 PWRD A paper recommended and approved by the IEEE Power System Instrumentation & Measurements Committee of the IEEE Power Engineering Society for presentation at the 1995 IEEE/PES Winter Meeting, January 29, to February 2, 1995, New York, NY. Manuscript submitted July 22, 1994; made available for printing November 2, 1994.

The nonintrusive load monitor (NILM) determines the operating schedule of the major electrical loads in a building from measurements made solely at the utility service entry [2], [3], [4] (See also [5] for an example of a related effort.). For electric utilities and industrial facilities managers, the NILM is a convenient and economical means of acquiring accurate energy usage data with a minimal installation effort. In [1], [6], and [7], a multi-scale transient event detection algorithm was introduced that can identify individual loads operating in a building by examining measured transient profiles observed in the aggregated current waveforms available at the service entry. This detection algorithm extends the applicability of the NILM to demanding commercial and industrial sites, where substantial efforts, e.g., power factor correction and load balancing, are made to homogenize the steady-state behavior of different loads, and where loads may turn on and off very frequently. With the incorporation of the transient event detector, the NILM is also an important platform for monitoring the performance and power quality of critical loads. For example, the NILM can track down power quality offenders, i.e., loads which draw extremely distorted, non-sinusoidal input current waveforms, by correlating the introduction of undesired harmonics with the operation of specific loads at a target site.

The next three sections of this paper describe, respectively, the theoretical foundation, implementation, and experimental results for an advanced preprocessor, or spectral envelope estimator, for use in the prototype event detector. The following three sections present the details of a transient event detection algorithm that identifies turn-on transients by examining spectral envelopes. The performance of the algorithm is demonstrated in the concluding sections with results taken from a prototype real-time event detector implemented with a digital signal processor.

II. Spectral Envelope Estimation

A critical observation made during the development of the prototype event detector is that direct examination

of a current waveform at the service entry, or a closely related waveform such as instantaneous power, may not satisfactorily reveal key features for transient identification. It is essential to distinguish such features from near constant frequency, "carrier wave" components of signals like 120 Hz instantaneous power so that slight errors in matching carrier frequency with a template do not dominate the results of a recognition system searching for a modulating envelope. A preprocessor in the prototype event detector eliminates carrier frequency artifacts from input data by averaging over at least one carrier wave period to generate a short time estimate of spectral content.

The development of the spectral envelope estimator is stimulated by the generalized averaging techniques presented in [8] and by the short time Fourier transform and Fourier series methods presented, for example, in [9] and [10] for speech processing and power systems simulation, respectively. With minor restrictions which cause no limitations in a practical power systems setting, a waveform $x(\tau)$ given as a function of τ may be described with arbitrary accuracy at time $\tau \in (t - T, t]$ by a Fourier series with *time-varying*, complex spectral coefficients $a_k(t)$ and $b_k(t)$:

$$x(t - T + s) = \sum_k a_k(t) \cos(k \frac{2\pi}{T}(t - T + s)) + \sum_k b_k(t) \sin(k \frac{2\pi}{T}(t - T + s)) \quad (1)$$

The variable k ranges over the set of non-negative integers; T is a real period of time, and $s \in (0, T]$.

The coefficients $a_k(t)$ may be found from the formula [8], [11]:

$$a_k(t) = \frac{2}{T} \int_0^T x(t - T + s) \cos(k \frac{2\pi}{T}(t - T + s)) ds \quad (2)$$

Similarly, the coefficients $b_k(t)$ are computed by the formula:

$$b_k(t) = \frac{2}{T} \int_0^T x(t - T + s) \sin(k \frac{2\pi}{T}(t - T + s)) ds \quad (3)$$

In practice, Eqns. 2 and 3 can be used to compute the evolution in time of the spectral coefficients $a_k(t)$ and $b_k(t)$ as an interval of interest of width T slides over the waveform x . The coefficients $a_k(t)$ and $b_k(t)$ as functions of time will be referred to as the *spectral envelopes* of x for the harmonic k .

Estimates of the spectral envelopes of current waveforms observed at the utility service entry of a building have proven remarkably useful for transient event detection in the NILM, for at least two reasons. First, even for waveforms x with substantial high frequency content, the frequency content of the spectral envelopes can be made relatively band-limited. As will be seen, this tends to ease the sample rate requirements on any

single channel of the transient event detector. Second, in steady-state operation especially, estimates of the spectral envelopes serve as direct indicators of real and reactive power, as well as potentially undesirable harmonic content. Demonstrations of these claims follow.

For convenience, let

$$x_c(t) = \frac{2}{T} x(t) \cos(k \frac{2\pi}{T} t)$$

and

$$x_s(t) = \frac{2}{T} x(t) \sin(k \frac{2\pi}{T} t).$$

represent sinusoids modulated by the function x . Equations 2 and 3 are equivalent to convolving in time the integrands $x_c(t)$ and $x_s(t)$, respectively, with a rectangular pulse $p(t)$ with unit height extending from time 0 to time T , and may be written as

$$a_k(t) = x_c(t) \otimes p(t) \quad (4)$$

and

$$b_k(t) = x_s(t) \otimes p(t) \quad (5)$$

where the symbol \otimes represents the convolution operator. In the frequency domain, the continuous time Fourier transform of the spectral envelope $a_k(t)$ is

$$A_k(f) = \int_{-\infty}^{\infty} a_k(t) e^{-j2\pi f t} dt.$$

From Eq. 4, the magnitude of $A_k(f)$ is equivalent to the product of the magnitudes of the functions $X_c(f)$, the Fourier transform of $x_c(t)$, and $P(f)$, the Fourier transform of the pulse $p(t)$. Similarly from Eq. 5, the magnitude of $B_k(f)$, the continuous time Fourier transform of $b_k(t)$, is the product of the magnitudes of $X_s(f)$, the transform of $x_s(t)$, and $P(f)$.

The effect of computing the spectral coefficients as an integral or average over the interval T is to attenuate the high frequency content of the spectral envelopes. Equivalently, the high frequency content of the spectral envelopes is attenuated by the (roughly) low-pass character of $P(f)$. The localization or high frequency attenuation in the frequency content of the spectral envelopes increases as the interval T increases in extent.

Each spectral coefficient indicates as a function of time the relative contribution of a sinusoid in the summations of Eq. 1. By varying the interval T it is possible to restrict to an essentially arbitrary degree the frequency content of the spectral envelopes, regardless of the harmonic k under consideration. In an actual implementation of a transient event detector, a decomposition of even a relatively broad-band waveform x into spectral envelopes permits a trade-off, therefore, between sample rate per data channel and the number of data channels employed.

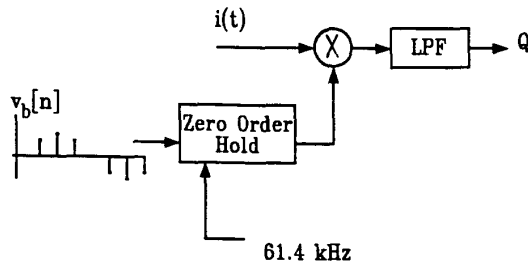


Figure 1: Signal Flow Graph

A second advantage of examining a waveform x in terms of spectral envelopes is the potential correspondence of the coefficients to familiar physical quantities in steady-state. The term “steady-state” is here taken to refer to a waveform or section of a waveform that is periodic. The interval T is presumed to be a positive integer multiple of the fundamental period of this waveform. Consider, for example, the situation in which the waveform x corresponds to an observed current waveform on a single phase of a utility system with a sinusoidal voltage waveform. For purposes of illustration, consider the voltage to be a cosine with angular frequency $2\pi/T$. Intuitively, Eqns. 2 and 3 compute the spectral coefficients by demodulating the periodic waveform x with an appropriate, harmonic sinusoid and low-pass filtering to preserve only the resulting lowest frequency components. For a periodic current waveform x with period T , the spectral coefficient $a_1(t)$ corresponds to a quantity that is proportional to the conventional definition of real or “time average” power [12]. Similarly, the coefficient $b_1(t)$ is proportional to reactive power. Higher order spectral coefficients correspond to in-phase and quadrature harmonic content as in a conventional Fourier series decomposition of a periodic waveform.

Without attempt at a mathematical exposition, we observe in passing that even for waveforms which are not strictly periodic, spectral envelopes may have appealing interpretations as quantities like real and reactive power, even though such quantities are not universally defined for transient or non-periodic waveforms. For waveforms that satisfy the “slowly varying magnitude and phase” arguments commonly used to justify quasi-sinusoidal steady-state approximations [10], the spectral envelopes can be loosely interpreted as the slowly varying envelopes of real and reactive power, and of harmonic content.

While not necessarily essential, the ability to associate spectral envelopes with physical quantities is comforting

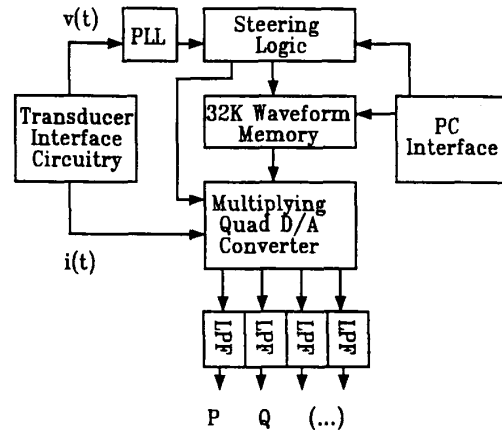


Figure 2: Envelope Preprocessor Block Diagram

when employing spectral envelopes as “fingerprints” in a transient event detector. Variations in real and reactive power and harmonic content tend to be closely linked to the physical task or energy conversion process being performed by a load. Load classes that perform physically different tasks are therefore distinguishable based on their transient behavior [1]. Since the spectral envelopes are closely linked to telltale physical quantities, they serve as reliable metrics for identifying loads.

III. Envelope Preprocessor

This section reviews the design of a hardware implementation of a spectral envelope preprocessor for use in the transient event detector. To maximize flexibility, utility, and accuracy, while minimizing cost, the preprocessor computes an estimate that approximates the spectral envelope integrals of Eqns. 2 and 3. Figure 1 illustrates the computation implemented in a single channel of the spectral envelope preprocessor. An observed current waveform $i(t)$, equivalent to the waveform x in the previous section, is mixed with a continuous time “staircase” or basis sinusoid. This basis sinusoid is constructed from discrete time samples, $v_b[n]$, of a desired waveform sampled at high frequency. Analytically, these samples are reconstituted into a continuous time waveform by a zero order hold (ZOH). This process is line-locked to the observed voltage waveform at the utility service entry, so that the reconstituted, basis sinusoid will exhibit a precise, desired phase with respect to the line voltage. The product of the current and basis sinusoid corresponds to a function such as $x_c(t)$ or $x_s(t)$ for a particular harmonic k , as described in the previous section. This product is low-pass filtered (LPF) to yield an estimate of a particular spectral envelope for the current.

Figure 2 shows a partial block diagram of the hardware implementation. A multiplying digital-to-analog converter (MDAC) provides the ZOH and multiplication operations shown in Fig. 1. The observed current waveform $i(t)$ is the analog reference for the MDAC. A memory on the preprocessor contains 1024 two byte samples of the desired basis waveforms for 16 different channels or spectral envelope estimates. These samples are fed to appropriate MDAC channels under the control of the steering logic. Although only four output channels are shown in Fig. 2, a total of 16 channels are available on the prototype envelope preprocessor. A phase-locked loop (PLL) ensures that the operation of the preprocessor is synchronized to the line voltage waveform. For enhanced flexibility, the basis sample memory may be implemented either with a read-only memory, or with a static RAM which can be loaded with samples from a personal computer.

In the prototype, a second order Butterworth filter with a breakpoint at 20 Hz is used on each channel to provide an estimate of the average or low-frequency component of each MDAC output. This low-pass filter, convenient from an implementation standpoint, is obviously not functionally identical to the windowed mean employed in Eqns. 2 and 3 to compute the spectral envelopes. For this reason, and because the basis waveforms are reconstructed with a (generally negligible) quantization error, the outputs of the prototype are estimates of the spectral envelopes. By varying the filter breakpoint, it is again possible to trade localization in time versus localization in frequency, as was possible in the previous section by varying the interval T .

IV. Envelope Preprocessor Performance

Three examples of the envelope estimator in operation will be considered here. In each case, the spectral envelopes will be observed to exhibit patterns created by the unique electrical, thermal, and mechanical processes that occur in each class of loads during the turn-on transient. Figure 3, for example, shows the measured current waveform and two associated spectral envelopes estimated by the preprocessor during the turn-on transient of a personal computer. All three waveforms were captured with a digital storage oscilloscope. The top trace in Fig. 3 shows the current into the computer during the transient. The switching power supply inside the computer initially draws a few large pulses of current as the internal bus capacitor charges from the line through a full bridge rectifier. When the capacitor has built up a substantial stored charge, the current waveform becomes "spikey" as charging begins to occur only near the peaks of the magnitude of the line voltage waveform. Approximately 0.12 seconds into the transient, the computer monitor turns on, increasing the total steady-state current drawn by the computer and monitor.




Figure 3: Personal Computer Turn-On Transient




Figure 4: Induction Motor Turn-On Transient

The second and third traces in Fig. 3 show two spectral envelope estimates computed by the preprocessor. The trace labeled a_1 corresponds to an envelope computed by mixing the observed current waveform with a sinusoid with the same phase and frequency as the line voltage. This trace corresponds to the slow envelope of "real power." The second trace, labeled a_3 , is computed by mixing the current with a sinusoid with the same phase but three times the frequency of the line voltage. This spectral envelope indicates in-phase third harmonic content. As might be intuitively expected from examining the "spikey" line current waveform, the computer draws a substantial third harmonic current.

Figure 4 shows the current and two spectral envelopes during the turn-on transient of a single phase induction




Figure 5: Compact Fluorescent Turn-On Transient

motor. While the rotor is accelerating to steady-state speed, the motor draws a relatively large current, as shown in the top trace of Fig. 4. Around 0.3 seconds into the transient, the rotor has reached its nominal velocity and the input current tapers to its steady state level. The trace labeled a_1 in Fig. 4 again corresponds to the slow envelope of real power. The spectral envelope preprocessor computed the third trace in Fig. 4, labeled b_1 , by mixing the current waveform with a sinusoid which is at the same frequency as, but which is ninety degrees out of phase with respect to, the line voltage waveform. This trace corresponds to the slow envelope of "reactive power." Again as might be expected, the induction motor does not exhibit unity power factor, and the magnitude of the b_1 trace in steady state is substantial relative to the magnitude of the a_1 trace.

Finally, Fig. 5 shows the current and three spectral envelopes measured during the turn-on transient of a bank of compact fluorescent lamps. Each lamp contains a ballast with an iron-core transformer and also power electronic circuitry. The combination draws a remarkably distorted input current waveform, including a significant in-rush current [13], as shown in the top trace in Fig. 5. The spectral envelopes corresponding to the slow envelopes of real power, reactive power, and in-phase third harmonic content are again labeled a_1 , b_1 , and a_3 respectively, and are shown below the current waveform. Even in steady-state, this load does not operate with unity power factor, as shown by the b_1 trace. Also, the "spikey" current waveform contains a substantial third harmonic component, as in the case of the personal computer. This is indicated by the a_3 trace in Fig. 5.




Figure 6: Measured Instant Start Lamp Bank Real Power Transient




Figure 7: Measured Induction Motor Real Power Transient

V. Approach to Transient Recognition

Because spectral envelopes are inherently band-limited and can be associated with familiar physical quantities related to the operation of different electrical loads, they have proven remarkably useful and robust for performing transient event detection [1], [6]. Nevertheless, searching for complete transients, even in spectral envelopes, is an undesirable approach because it limits the tolerable rate of event generation. No two transient events could overlap significantly if each transient were to be identified correctly. Instead, the transient event detector searches for a time pattern of segments with significant variation, or *v-sections*, rather than searching for a transient shape in its entirety.

During a training phase, either before installation or on-site, the event detector employs a change-of-mean detector to segment a transient representative of a class of loads [14]. This segmentation process delineates a set of *v-sections* that will represent a particular transient shape in each of the input spectral envelopes. The trace in Fig. 6, for example, shows the measured envelope of real




Figure 8: Acceptable Overlap




Figure 9: Intractable Overlap

power during the turn-on transient of an instant start fluorescent lamp bank. Figure 7 shows the measured envelope of real power on one phase during the turn-on transient of a three phase induction motor. The locations of the v-sections in the two waveforms, as computed by a change-of-mean detector implemented in MATLAB, are approximately indicated by the ellipse in Fig. 6 and the rectangles in Fig. 7. In practice, a more complicated set which included v-sections in other spectral envelopes, i.e., estimates of reactive power and higher harmonic content, would be used to represent the transient profile of a load.

A complete transient identification is made by searching for a precise time pattern of v-sections. As long as each of the v-section shapes overlaps with no more than a near-constant region, the event detector will be able to identify the patterns of v-sections and therefore the transients. For example, the overlap of the two transients from the induction motor and the instant start lamp bank shown in Fig. 8 is tractable because all of the v-sections for both transients are distinguishable. The overlap condition in Fig. 9 would not generally be tractable, since the instant start v-section and the first

induction motor v-section overlap severely. Since some degree of overlap is tolerable, the v-section set recognition technique will generally operate successfully in an environment with a higher rate of event generation than would a detector searching for whole, undisturbed transient shapes.

Classification of individual v-sections in the input data streams is determined by a set of pattern discriminator functions. These functions are used to compute a distance metric that locates a particular input vector in a region of a state space of known transient templates. Since a v-section may appear on top of a variably large static or quasi-static level created by the operation of other loads, the discrimination process focuses on only the "AC" or varying component of the v-section.

The prototype event detector employs a transversal or matched filter as a pattern discriminator, although other possibilities could be used and are discussed in [1]. The impulse response of a transversal filter is proportional to the time-reversed signal for which the filter is designed to search. The transversal filter is an attractive signal processing construct for performing pattern discrimination because there are many efficient hardware solutions available for implementing a transversal filter. Each v-section is positively identified by checking the outputs of two different transversal filters.

The first transversal filter scans an input data stream for a particular shape. Let t denote a vector of N samples of a v-section of interest, such as one of the v-sections marked in Figs. 6 and 7. The vector t consists of elements

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

An "AC coupled" and amplitude normalized version of the v-section, designated t_{ac} , can be computed as

$$t_{ac} = \frac{\left[t - \frac{1}{N} \sum_{i=0}^{N-1} t[i] \right]}{\left\| \left[t - \frac{1}{N} \sum_{i=0}^{N-1} t[i] \right] \right\|} \quad (6)$$

where the denominator in Eq. 6 is the 2-norm of the numerator. Thus, t_{ac} is a unit length vector with zero mean. This shape transversal filter operates by sliding an N -point window across an input spectral envelope. At any time, the window contains an AC-coupled and amplitude normalized vector of points x_{ac} of an N -point section or vector x of the input data. The vector x_{ac} is computed from the vector x using an equation structurally identical to Eq. 6. The output of the shape transversal filter is the inner product of the template vector t_{ac} and the data vector x_{ac} .

This output corresponds to the cosine of the angle between the two unit length vectors t_{ac} and x_{ac} . An output of unity indicates a perfect match between the template vector and the input data. Naturally, noise and slight variation in the repeatability of the v-sections will make

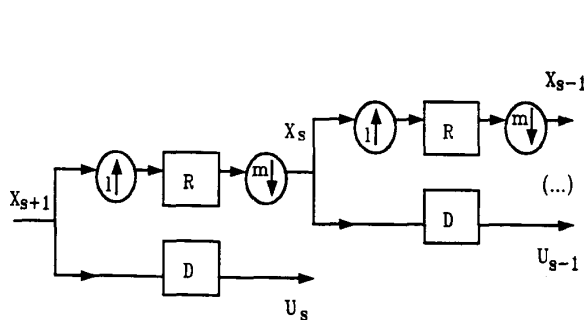


Figure 10: Tree-Structured Signal Decomposition

a perfect match unlikely. In a practical system some degree of imperfection will be tolerated and any inner product within a certain tolerance of unity will constitute a match.

When a segment in the input data stream is found that matches the shape of a template v-section, a second transversal filter is employed to check the amplitude of the segment. This amplitude transversal filter computes the inner product of the template vector t_{ac} and the input data vector x , rather than the inner product of t_{ac} and x_{ac} . Checking the amplitude is essential in conjunction with checking shape to ensure that a small wiggle or noise pattern that is fortuitously close in shape to a v-section template is not mistaken for an actual v-section.

VI. Searching for Patterns Over Many Time Scales

Loads in a particular class which spans a wide power range often exhibit transient profiles that are identical in shape but scaled in amplitude and duration. The transversal filter is suitable for identifying transient shapes over a narrowly defined time scale. The prototype event detector employs a tree-structured decomposition to search efficiently over many time scales with the transversal filters. The use of the tree-structured decomposition is inspired by recent applications of subband coding [15] and the discrete-time wavelet transform [16], [17].

A tree-structured signal decomposition is shown schematically in Fig. 10. An atomic section, or coder, of the tree shown would consist of the upsampler with upsampling rate l , operations R and D , and the trailing downsampler with downsampling rate m . Each upper path, or *resolving path*, in any particular coder in the tree alters the resolution of the input signal with operation R , and also alters the scale of the input signal with the up/downsampling operations at rates l and m . The lower *discriminating path* implements operation D which

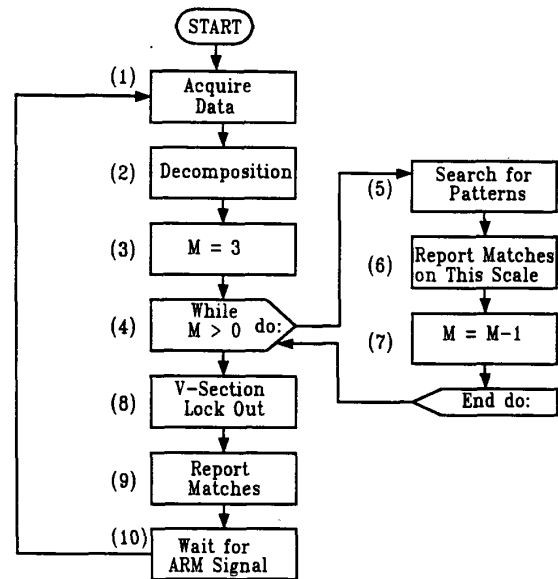


Figure 11: The NILMscope Event Detection Algorithm

identifies patterns at a particular time scale in the input data. In the prototype event detector, a transversal filter computes the discriminating coefficients, U_s , for example, that indicate the possible presence of a v-section.

Experience gathered from the load survey in [1] and from experiments with the event detector indicate that operation R may be implemented with a linear filter for the v-sections associated with many commercial and industrial loads. Constraints for designing filters that minimize the impact of distortions in a pattern recognition setting, and conditions when nonlinear filters might be desirable, are presented in [1].

VII. Event Detection Algorithm

A flow chart of the complete event detection algorithm is shown in Fig. 11. This algorithm is implemented in the *NILMscope* software that is executed by the prototype transient detector. The *NILMscope* software is written in two components. One component performs the event detection algorithm and is executed on a TMS320C30 DSP system. The other is a user interface responsible for initializing the DSP system and for report generation. This second component is executed on an 80486-based personal computer. Each numbered item in the following list is a description that details a numbered step of the algorithm in the flow chart:

1. Data Acquisition: During this data acquisition step, the DSP system collects a window of samples which will be searched for known transient patterns.

2. **Tree-Structured Decomposition:** Once a full window or vector of samples has been acquired, the DSP system performs a tree-structured decomposition on the data as described in the previous section. In the current implementation, the input data for each coder step in the tree is computed before any pattern discrimination occurs at any scale step. A tree structure with a total of three 2 to 1 coder or scale steps proved sufficient for identifying all of the transients associated with the loads in the test stand described in the next section.
3. **Set Scale Steps:** Next, the DSP system searches at each scale for all of the transient types that could appear. There are three scale steps, and a scale step variable M is initialized to the value of 3 to count down the search. First the finest sampled scale step is inspected, followed by the middle and coarse scales.
4. **Initiate Pattern Search:** A loop in the program flow at this point in the *NILMScope* software controls the search for patterns over all three scale steps.
5. **Hierarchical Pattern Search with V-Section Lock Out on Scale M :** During each pass through the pattern search loop, the DSP system searches for the v-sections associated with the known transient events on a single scale. The pattern search is hierarchical, in that the DSP system searches first for the patterns with the most v-sections. When all of the v-sections for a pattern are found with both the shape and amplitude transversal filtering operations, the complete transient pattern is presumed to be present in the input data, and an event is recorded.

A v-section lock out is performed at each scale. If a complex pattern is found in the input data, the location of the v-sections of the pattern are recorded. The identification of any subsequent, less complex patterns will not be permitted based on the detection of v-sections at the previously recorded, "locked out" locations.
6. **Report Generation for Scale M :** If all of the v-sections are found for a particular pattern, the transient pattern is presumed to be present in the data at the current scale M , and an event type and time is communicated to the interface component of the *NILMScope* software running on the PC.
7. **Decrement Scale Step:** The scale counter M is decremented, and the pattern detection loop is repeated until all remaining coarser scales have been searched.
8. **V-Section Lock Out Over All Scales:** The PC component of the *NILMScope* software performs a final check to ensure that v-sections from a complex but

coarse scale pattern were not used to match a less complicated, finer scale pattern.

9. **Final Report Generation:** The PC component of the *NILMScope* software collates a final report of the type and time of occurrence of all positive event detections.
10. **Standby:** After reporting any contacts, the PC waits for the user to issue an arming command.

VIII. Prototype Testing and Performance

The prototype event detector consists of three components: the spectral envelope preprocessor, a DSP card, and a personal computer. The event detector monitors the voltage and current waveforms on a three phase electrical service that powers a collection of loads representative of important load classes in typical, medium to large size commercial and industrial buildings. The prototype event detector is used to identify the turn-on time and type of the various loads. The event detector has, of course, no a priori knowledge of the operating schedule for the loads.

Four loads were selected for inclusion in the test facility: a bank of instant start fluorescent lamps, a bank of rapid start fluorescent lamps, a $\frac{1}{4}$ horsepower, three phase induction motor, and a $\frac{1}{3}$ horsepower, three phase induction motor. The electrical hookup to the loads is routed through an electronically switched circuit breaker panel that activates loads with flexibility in relative timing. The pattern templates for the loads were captured during a one-time "walk through" of the test stand. However, no data at all was collected from the large motor. Because the large and small induction motors are members of the same load class, a single transient template, appropriately scaled in amplitude and duration, was expected to prove satisfactory for identifying both motors.

Figures 12 through 17 show screen prints from the PC running the *NILMScope* user interface software during six of the experiments conducted with the test stand. The graph windows in each figure show estimates of the envelopes of real and reactive power on one phase of the three phase service. For example, the windows in Fig. 12 display the data collected during the turn-on of the instant start lamp bank in the test stand.

In the lower left-hand corner of each screen, any transient events that the event detector has been able to identify appear under the heading **Contacts**. Events or contacts are reported by identity, time of occurrence, and scale. There are three scales listed in the contact window: fine, mid, and coarse. Any event that was identified by a known v-section set at the initial, highest sampling rate, i.e., at the first coder stage in the tree-structured decomposition, will be listed directly under the heading




Figure 12: NILMScope Contact Report

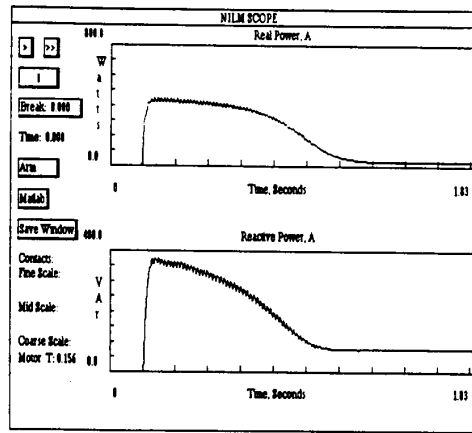


Figure 15: NILMScope Contact Report




Figure 13: NILMScope Contact Report




Figure 16: NILMScope Contact Report




Figure 14: NILMScope Contact Report




Figure 17: NILMScope Contact Report

Fine Scale in the contact window. By design, it is anticipated that events associated with the small motor and both lamp banks will be listed as fine scale events when they appear, assuming that the event detector functions properly. Events found at the next two coder stages in the tree-structured decomposition will be listed under the headings **Mid Scale** or **Coarse Scale**. For example, any events recognized by the properly working detector which are caused by the turn-on of the large induction motor should appear under the coarse heading.

The four tests shown in Figs. 12 through 15 record the performance of the prototype when challenged individually with the turn-on events of the instant and rapid start lamp banks and the small and large induction motors, respectively. In each case, the observed event has been properly identified in the contact window. Figure 16 shows an example where both lamp banks and the small motor turn on so that all three transient events overlap. No key v-sections overlap with each other. All three events are correctly recorded at the finest time scale in the contact window.

In the final experiment shown in Fig. 17, the small induction motor turns on and completes its transient, followed by the turn-on transient of the large induction motor and the instant start lamp bank. No key v-sections overlap with each other. All of the events are correctly identified at the appropriate time scales. Recall that the template for the turn-on transient of event type **Motor** was generated from a single example of the small motor only. Nevertheless, the event detector correctly classified both the small and large motors in the experiments shown in Figs. 14 through 17.

IX. Conclusions

The examples reviewed in the previous section are representative of several hundred experiments conducted with the test stand. Provided the assumptions made in the development of the event detection algorithm are satisfied, the prototype detector performs remarkably well. This performance is perhaps more impressive in light of the fact that fairly little effort was made to "tune" the detector for the loads in the test stand.

In practice, the robustness of the detector could be enhanced by working with more complicated v-section pattern sets for each load. Employing more complicated sets of interesting v-sections will enhance the reliability of the event detection algorithm but will also require more computational capability. Fortunately, the search for different v-sections could be conducted in parallel. This suggests that a commercial NILM based on the event detection algorithm could be constructed around a parallel processing architecture composed of inexpensive microprocessors or microcontrollers.

Acknowledgements

This research was funded by EPRI and ESEERCo. The authors gratefully acknowledge the valuable advice and support of Professors George Verghese, Chris DeMarco, Bernie Lesieutre, and Les Norford, Dr. Richard Tabors, Frank Porretto, Laurence Carmichael, and Kurt Levens. Essential hardware for this project was made available through a generous donation from the Intel Corporation. In particular, the authors wish to thank Ann Bynum of Intel for her tireless help and interest in the Laboratory for Electromagnetic and Electronic Systems at MIT.

References

- [1] Leeb, S.B., "A Conjoint Pattern Recognition Approach to Nonintrusive Load Monitoring," M.I.T. Ph.D. Dissertation, Department of Electrical Engineering and Computer Science, February 1993.
- [2] Tabors, R., et. al., "Residential Nonintrusive Appliance Load Monitor," EPRI final report, to appear.
- [3] Kern, E.C., "Meter Sorts Load for Appliance Use," *Electrical World*, p. 75, April 1986.
- [4] Hart, G.W., "Nonintrusive Appliance Load Monitoring," *Proceedings of the IEEE*, Vol. 80, No. 12, pp. 1870-1891, December 1992.
- [5] Sultanem, F., "Using Appliance Signatures for Monitoring Residential Loads at Meter Panel Level," *IEEE PES Winter Meeting*, No. 91 WM 016-6 PWRD, New York, New York, February 1991.
- [6] Leeb, S.B., and J.L. Kirtley, Jr., "A Multiresolution Transient Event Detector for Nonintrusive Load Monitoring," *International Conference on Industrial Electronics, Control, and Instrumentation 1993*, Maui, Hawaii, pp. 354-359, November 1993.
- [7] Leeb, S.B., and J.L. Kirtley, Jr., "A Transient Event Detector for Nonintrusive Load Monitoring," U.S. Patent Application, filed December 1993.
- [8] Sanders, S.R., J.M. Noworolski, X.Z. Liu, and G.C. Verghese, "Generalized Averaging Method for Power Conversion Circuits," *IEEE Transactions on Power Electronic Circuits*, Vol. 6, No.2, pp. 251-259, April 1991.
- [9] Oppenheim, A.V., *Applications of Digital Signal Processing*, Prentice-Hall, Inc., Engelwood Cliffs, New Jersey, pp. 117-168.
- [10] DeMarco, C.L., and G.C. Verghese, "Bringing Phasor Dynamics into the Power System Load Flow," *North American Power Symposium*, 1993.
- [11] Siebert, W.M., *Circuits, Signals, and Systems*, McGraw-Hill, New York, New York, p. 370.
- [12] Kirtley, Jr., J.L., "AC Power Flow in Linear Networks," M.I.T. 6.061 Supplementary Notes 2, June 1987.

- [13] Slemon, G.R., *Magnetolectric Devices: Transducers, Transformers, and Machines*, John-Wiley and Sons, 1966.
- [14] Basseville, M., and A. Benveniste, *Detection of Abrupt Changes in Signals and Dynamical Systems*, Springer-Verlag, 1980.
- [15] Vaidyanathan, P.P., "Multirate Digital Filters, Filter Banks, Polyphase Networks, and Applications: A Tutorial," *Proceedings of the IEEE*, Vol. 78, No. 1, January 1990, pp. 56-93.
- [16] Mallat, S.G., "A Theory for Multiresolution Signal Decomposition: The Wavelet Representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 11, No. 7, July 1989.
- [17] Meyer, Y., "Orthonormal Wavelets," *Wavelets: Time-Frequency Methods and Phase Space*, Proceedings of the International Conference, Marseille, France, December 14-18, 1987, Springer-Verlag, pp. 21-37.

Steven B. Leeb (S'89-M'93) received his Bachelor of Science, Master of Science, Electrical Engineer, and Doctoral degrees from the Massachusetts Institute of Technology in 1987, 1989, 1990, and 1993, respectively. He has been a member of the M.I.T. faculty in the Department of Electrical Engineering and Computer Science since 1993. His appointment to the faculty is as an Assistant Professor in the Laboratory for Electromagnetic and Electronic Systems.

Dr. Leeb is concerned with the design, analysis, development, and maintenance processes for all kinds of machinery with electrical actuators, sensors, or power electronic drives. He is a member of the IEEE Power Electronics, Control Systems, Power Engineering, and Signal Processing Societies. A Fellow of M.I.T.'s Leader's for Manufacturing Program, he is also a member of Tau Beta Pi and Eta Kappa Nu.

Steven R. Shaw is an undergraduate research assistant in the Laboratory for Electromagnetic and Electronic Systems at MIT. He is a senior in the Department of Electrical Engineering and Computer Science.

James L. Kirtley, Jr. (S'69-M'71-SM'80-F'91) received the S.B. and Ph.D. degrees from the Massachusetts Institute of Technology, Cambridge, in 1968 and 1971, respectively.

He joined the faculty in 1971 and is now Professor of Electrical Engineering at MIT. Dr. Kirtley is Vice Chairman and Technical Paper Coordinator of the Electric Machinery Committee, and a member of the Synchronous Machinery Subcommittee and the Machinery Theory Subcommittee. He was chairman of the 1990 International Conference on Electric Machines. He is a member of CIGRE, a member of the editorial board of *Electric Machines and Power Systems*, and a Registered Professional Engineer in Massachusetts.