

S. R. Shaw

C. B. Abler

R. F. Lepard

D. Luo

S. B. Leeb

L. K. Norford

Massachusetts Institute of Technology
LEES Laboratory
Cambridge, MA 02139

Instrumentation for High Performance Nonintrusive Electrical Load Monitoring

This paper reviews the design and implementation of hardware and software tools for nonintrusive electrical load monitoring. Estimates of spectral content in measured waveforms can be used to determine in real time the operating schedule of loads at a target site. Techniques for transient event detection are reviewed. These techniques can detect the turn-on and turn-off transients of individual loads, and can be used to easily determine the energy usage of loads that draw constant power in steady-state operation. Techniques for monitoring the power consumption of smoothly varying loads (e.g., variable speed drives) using spectral estimates are also discussed.

I Load Monitoring

The availability of accurate, inexpensive usage monitoring for electrical loads makes it possible for commercial and industrial facilities managers to minimize operating costs and downtime for repairs. Detailed information about the operation of electrical loads can also be used for power quality monitoring and the diagnostic determination of the health of critical loads, e.g., lighting, HVAC components, and process-line equipment in manufacturing operations. Electric utilities and service providers want accurate usage information to refine load models used in security assessments and stability analyses for distribution systems. Energy service companies and facility owners would like an inexpensive means to measure and verify savings from energy efficiency improvements. The commercial market has shown dramatic advances in the availability, sophistication, and affordability of monitoring instrumentation in the last two decades.

This paper reviews our recent work to develop signal processing algorithms and hardware for nonintrusive load monitoring. The nonintrusive load monitor (NILM) is a device that can determine the operating schedule of all of the significant electrical loads at a target site strictly from voltage and current measurements made at the electric utility service entry [5], [3], [6], [8], [7], [13]. The minimal effort associated with installing a NILM in comparison to traditional submetering, as well as the convenience of retrieving collated data from a single monitoring device, make the NILM a very attractive load monitoring platform for many different applications. The earliest approach to nonintrusive monitoring determined the operating schedule of loads of interest by detecting changes in the "steady-state" real- and reactive-power consumption at the service entry. This approach, still in use in commercially available NILMs, has difficulty differentiating loads in demanding commercial and industrial environments. Difficulties can arise when substantial efforts have been made to homogenize the steady-state energy consumption of loads through load balancing and power-factor correction, when loads turn on and off frequently, and when loads are present that have a continuously variable power consumption, e.g., a variable-speed drive (VSD).

Transient event detection can be used to successfully monitor sites with complicated loads and combinations of loads [11], [4]. As loads turn on, the current waveform observed by the NILM will contain transient shapes. The turn-on transients associated with many load classes are sufficiently distinct and repeatable to associate transient observations with the operation of particular loads. However, we have found that direct examination of the current waveform or a closely associated waveform like instantaneous power may fail to accentuate important features for pattern recognition. A recognition system searching for gross features may be overwhelmed by attempting to match the essentially irrelevant fine-structure of a current waveform. Time-varying estimates of the frequency content of the current at harmonics synchronized to the voltage have proven invaluable for load monitoring [10], [11], [9]. These spectral envelopes are closely associated with the physical tasks performed by a load. Also, the frequency content of any particular spectral envelope is relatively band-limited in comparison to the original, observed waveform. Hence, spectral envelopes enable a flexible trade-off between sample-rate requirements for any one data channel and the total number of channels required to adequately characterize a waveform. When it is desirable to perform NILM computations on a parallel processing machine, spectral envelopes also facilitate the distribution of data to individual processors [4].

Typically, energy consumption is estimated by assuming a near-constant steady-state power consumption for each load over an interval bounded by a turn-on and a turn-off transient. Even with transient event detection, therefore, the determination of the power consumption of a continuously variable load like a VSD has proven difficult. We have used spectral envelopes to develop one method for tracking the power consumption of some continuously variable loads with a NILM. Previous field tests of the transient event detector were conducted with a custom spectral envelope preprocessor [11]. Aggressive hardware performance advances have made it possible to implement a transient event detector entirely in software on a commercially available signal-processing computer. This paper reviews the mathematical techniques and software modules used to create a transient event detector on an inexpensive digital signal processor (DSP). The next two sections describe the calculation of spectral envelope estimates and the implementation of these calculations on a digital signal processor. The following sections describe techniques for using spectral envelopes to identify the operating schedules of both loads that draw constant steady-state power and also some continuously variable loads.

Contributed by the Solar Energy Division of The American Society of Mechanical Engineers for publication in the ASME JOURNAL OF SOLAR ENERGY ENGINEERING. Manuscript received by the ASME Solar Energy Division, Mar. 1998; final revision, May 1998. Associate Technical Editor: D. E. Claridge.

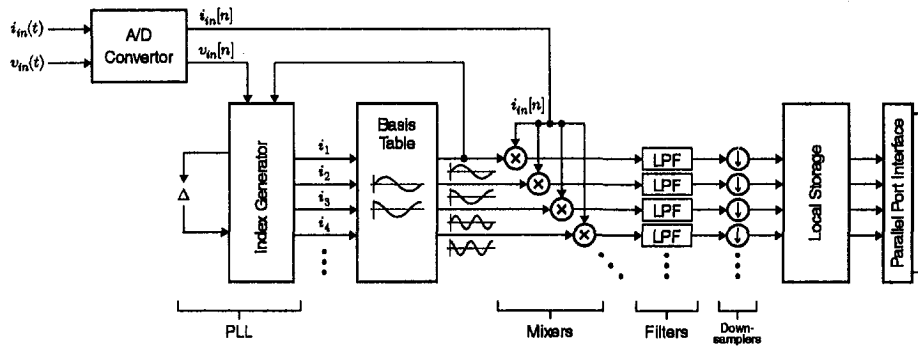


Fig. 1 Functional block diagram of the envelope estimator

II Spectral Envelope Estimation

Spectral envelopes are short-time averages corresponding to the time-local harmonic content of a waveform. For transient event detection on the AC utility, a spectral envelope is computed by integrating the product of the measured current waveform with a basis function line-locked to the measured voltage waveform over some finite interval. The basis functions are typically sinusoids at harmonics and phases of interest. For example, assume that the measured voltage waveform is a sine-wave at an angular frequency ω . The in-phase spectral envelopes a_k of the waveform x are then

$$a_k(t) = \frac{2}{T} \int_{t-T}^t x(\tau) \sin(k\omega\tau) d\tau, \quad (1)$$

where k is a non-negative integer that represents the harmonic index. Similarly, the quadrature spectral envelopes b_k are

$$b_k(t) = \frac{2}{T} \int_{t-T}^t x(\tau) \cos(k\omega\tau) d\tau. \quad (2)$$

The averaging interval T is typically one or more periods of the fundamental frequency of the voltage waveform. These spectral envelopes are the coefficients of a time-varying Fourier series of the waveform $x(t)$ [11]. With the basis functions in (1) and (2), note that spectral envelopes a_1 and b_1 have the same form as conventional definitions of real and reactive power in steady state.

Equations (1) and (2) can be interpreted as the convolution of the product of $x(t)$ and a basis sinusoid or cosinusoid with a pulse of unit height and time duration T . The pulse has a frequency response that is low-pass in character [12]. In practice, estimates of the spectral envelopes associated with $x(t)$ can be conveniently computed by low-pass filtering the product of $x(t)$ and the basis waveforms. Varying the cutoff frequency of the filter trades localization in time with localization in frequency. That is, a filter with a sharp, low-frequency cutoff point will clearly indicate the harmonic content of a narrow frequency range of $x(t)$. The spectral envelope estimate produced by this filter will be relatively band-limited and, therefore, easily sampled at a low rate. However, a low cutoff filter will respond slowly to changes in harmonic content. Alternatively, a low-pass filter with a broader pass band in frequency will require a higher sample rate in time, and will also respond relatively quickly to changes in harmonic content.

III Preprocessor Implementation

We have implemented a spectral envelope preprocessor to compute estimates of (1) and (2) in real time on an inexpensive, high performance digital signal processor (DSP). We chose Motorola's DSP56303, which is a modern fixed-point, 24-bit, 80 MIP DSP. The DSP evaluation board provides support hardware, including a sixteen bit two-channel analog-to-digital con-

verter, non-volatile and RAM memory, and a serial debugging interface. The preprocessor implements the block diagram shown in Fig. 1. The signal processing operations in Fig. 1 are conventional except for the software-emulated phase-locked loop (PLL), which adaptively indexes a basis function array to synchronize basis functions to the line. A flow chart of the code running on the DSP appears in Fig. 2. Samples of the input are collected continuously by interrupt and buffered while this flowchart executes. The "stop" condition, where the DSP simply idles and updates the PLL, can be invoked by the host computer's device driver when a NILM application is not running. A brief description of the key software operations follows; further details can be found in [1].

A Phase-Locked Loop. The software PLL design parallels the components of a hardware PLL. The software phase-locked loop computes a phase error, filters it, and adjusts the

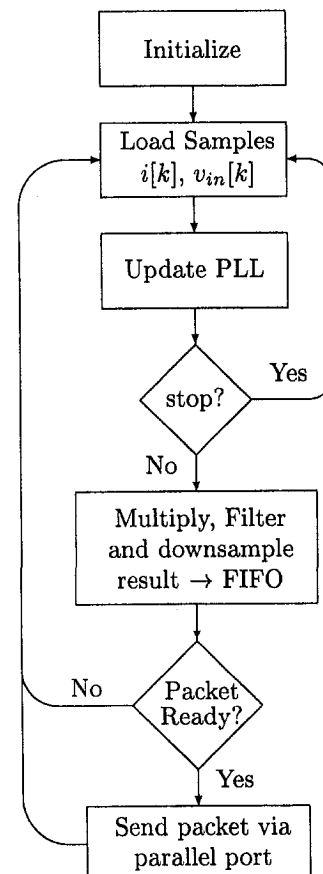


Fig. 2 Flow chart of spectral envelope estimation code




Fig. 3 State machine showing evolution of increment $\Delta[k]$

rate at which an index traverses a table containing a sampled sine wave. These operations are analogous to the phase-comparator, loop filter, and voltage-controlled oscillator in a conventional PLL circuit.

The phase error is computed by comparison of the zero crossings of two signals; the sampled voltage $v_{in}[k]$, and the reference voltage $v_{ref}[k]$. We create two signals r_{in} and r_{ref} indicating rising zero-crossing on the sampled voltage and reference waveforms, respectively. In particular,

$$r_{in} = E(v_{in}) \quad (3)$$

where the operation, $E(x) \mapsto y$ is given by the state machine in Fig. 4, and r_{ref} is similarly defined. The machine in Fig. 4 is designed to ameliorate the effects of noise at the zero crossing. In turn, r_{in} and r_{ref} drive a state machine that determines the evolution of the increment $\Delta[k]$ used to generate indices for the basis waveforms. An illustrative state machine is shown in Fig. 3; the state machine used in our prototype is similar but interpolates to locate zero crossings between samples [1]. The filter values α , β , etc., in Fig. 3 are determined from standard analog PLL loop-gain calculations, mapped to discrete time. The increment $\Delta[k]$ determines the evolution of the indices i_n of the basis waveforms,

$$i_n[k+1] = (i_n[k] + n\Delta[k]) \bmod N \quad (4)$$

where N is the length of the stored cosine and sine arrays C and S . For example, the "VCO output," $v_{ref}[k]$, is indexed by $i_1[k]$, i.e.

$$v_{ref}[k] = S[i_1[k]]. \quad (5)$$

Similarly, the third-harmonic cosine basis waveform at time k is $C[i_3[k]]$.

B Filtering and Data Transmission. The spectral envelope estimation process is completed by multiplying the measured current with the phase-locked basis waveforms and low-pass filtering. The results of the multiplication stage are accumulated in first-in, first-out buffers and low-pass filtered in blocks with a 256-tap finite-impulse response (FIR) filter. The FIR filters are attractive in pattern recognition applications because




Fig. 4 State machine defining the operator $E(x) \mapsto y$

of their linear-phase characteristic. In our system, the breakpoint of the filter is 20 Hz, band-limiting the spectral envelope estimates. The input sample rate of the system, i.e., the sample rate of the current and voltage waveforms, is 9600 Hz, ten times greater than the fastest significant frequency component of the measured current. The band-limited spectral envelope estimates are sent to the host computer at a sampling rate of 200 Hz, in blocks of 64 bytes, over a standard bi-directional parallel port interface. Deep buffers are implemented both on the DSP side of the link and on the PC side. On Intel Pentium-class systems running Linux, overhead for continuously receiving spectral envelope estimates is negligible.

C Results. For under \$200, the DSP evaluation board can be used to implement a highly effective spectral envelope preprocessor. The characteristics of the envelope calculation, for example, the FIR filter breakpoint, can quickly and easily be changed or tuned for a specific application. With eight channels, we estimate that the DSP chip is running at about twenty-five percent load; extra cycles could be used for additional analysis, or to calculate more spectral envelopes. In a field environment, the DSP envelope estimator has proven easily capable of detecting one watt changes on top of a 10,000 watt base load, i.e., a dynamic range of four orders of magnitude. Figure 5, for example, shows a spectral envelope (which corresponds to real power in steady state) during the turn-on transient of an incandescent light bulb. After a significant in-rush current when the bulb filament is cool, the power consumption settles to its nominal level. This measurement was made by a spectral envelope preprocessor with sensors configured to measure a 55 kW electric service.

IV Event Detection

Rather than searching for entire transients, the event detector in the NILM searches for segments of the spectral envelopes of an electrical transient with significant derivative or variation. These regions are called v -sections. During a training phase,




Fig. 5 Real power spectral envelope of a 25 W, incandescent light bulb

either before installation or on-site, the NILM employs a change-of-mean detector [2] or other waveform segmentation technique to identify a collection of v -sections representative of a load or class of loads. A complete transient identification is made by searching for a precise time pattern of v -sections in all available and informative spectral envelopes [11].

The v -sections are relatively narrow in duration, and our field experience indicates that they are less likely to be fatally corrupted by overlap than would an entire transient. The transient event detector searches for v -sections with an AC-coupled pattern discriminator such as a transversal filter [11] or Euclidean-distance filter [4]. Figure 6, for example, shows the real-power spectral envelope during the nearly simultaneous turn-on of an induction motor and an instant-start fluorescent lamp bank. The instant-start lamp transient, shown in the oval in Fig. 6, starts during the middle of the motor's acceleration to a steady-state operating speed. As long as each v -section overlaps with only quasi-static or nearly "flat" sections of other transients or steady-state operating curves, AC-coupled pattern discriminators can identify each v -section and successfully identify a transient. Since some degree of overlap is tolerable, the v -section set recognition technique will generally operate successfully in a busier environment with a higher rate of event generation than would a detector searching for whole, undisturbed transient shapes. In situations where two v -sections directly overlap, a decision must be made whether or not to search combinations of known v -sections to detect a match, risking a possible false identification.

A computationally efficient technique for detecting changes or transients in a signal is valuable in this setting for at least two reasons. First, such a detector can be used to segment a whole transient into v -sections while training the NILM, as discussed above. Second, the detector can be used to minimize the burden of searching spectral envelope estimates for v -sections by identifying only regions of the envelopes that show significant change, i.e., the possibility of the presence of a v -section of interest. This can free processing time on the NILM platform for "value-added" services such as diagnostic monitoring of critical loads [10].

V Change-of-Mean Detection

Change-of-mean detection serves two roles in nonintrusive electrical load monitoring. First, as noted, such detection usefully delimits transients, for training a load monitor to recognize certain patterns or v -sections and later, in on-line service, to flag power changes for transient analysis. Second, change-of-mean detectors can operate alone on the power envelope in environments where equipment operation is controlled, as with a building energy management system (BEMS). In this case, the identity of a load like a pump or chiller is known, the start-




Fig. 6 Turn-on transients: induction motor and fluorescent lamp




Fig. 7 Minute average NILM power data for building E23

up sequence typically is staggered such that power transients do not overlap, and the full power of v -section analysis is not required to estimate start-up or shut-down power associated with a particular piece of equipment.

Most change-detection algorithms are based on a ratio of probability density functions of sampled power data relative to mean levels μ_0 and μ_1 before and after an event [2], [14]. For example, the decision function S_j^k ,

$$S_j^k(\mu_1) = \sum_{i=j}^k \ln \frac{P(y[i] - \mu_1)}{P(y[i] - \mu_0)}, \quad (6)$$

sums the log-likelihood ratio over a sliding window of $k - j$ samples in some measurements $y[i]$, given a probability density function P . The generalized-likelihood ratio (GLR) test seeks to maximize the decision function by searching over the change time and the magnitude of the change in mean. For an independent Gaussian sequence with variance σ and a minimum change in mean power as low as zero, the maximization of the decision statistic over the sample time becomes

$$g_k = \frac{1}{2\sigma^2} \max_{1 \leq j \leq k} \frac{1}{k - j + 1} \left[\sum_{i=j}^k (y[i] - \mu_0) \right]^2. \quad (7)$$

A change in power due to the start-up or shut-down of a lamp bank or an HVAC component will be detected if g_k exceeds a specified threshold.

The pre-event sample mean μ_0 can be determined by a sliding-window average. The window must be sufficiently long to be stable in a noisy environment and not react too quickly to the onset of a gradual rather than abrupt start-up event; however, a long window causes the decision function to miss rapidly occurring events. To enhance sensitivity and robustness, we update the window with a new set of data and reset the decision statistic immediately after the detection threshold has been exceeded. The sample variance will vary with the mix of equipment in operation at any time; in accordance with (7), the decision statistic will vary as well. A fixed variance will lead to false alarms if that variance is low relative to a period of high sample noise and will cause events to be missed if the variance is larger than characterizes a period of little noise. To avoid this problem, we update the variance from the pre-event averaging window. The decision statistic can be readily modified to account for minimum expected changes of mean associated with known devices, as would be appropriate were the detector to be married to a BEMS. Such modification reduces false alarms that arise from detecting power changes less than those of interest. Figure 7 shows an electrical signal that is the




Fig. 8 GLR test for building E23 with recursive window reset

aggregate of one-minute average power data for eight variable-air-volume ventilation fans, each equipped with a variable-speed drive. Figure 8 gives the decision statistic, which detects all of the on-off events and two cases, near 200 and 1500 minutes, when a fan slows down rapidly, but fails to detect an abrupt slow-down in fan speed between two start-ups near 1200 minutes.

VI Variable Speed Drives

Variable-speed drives are important electrical loads that are difficult to monitor nonintrusively. They are important because they are used increasingly in the HVAC systems of many buildings to improve energy efficiency. Variable drives are often incorporated in servo-mechanical feedback loops in which the drive speed is varied to regulate a pressure or flow velocity in the presence of disturbances. Also, the set point or reference for the controlled variable may vary smoothly throughout the course of an operating cycle. Therefore, the net energy consumption of a VSD cannot generally be computed by assuming a constant power consumption over a time interval bounded by turn-on and turn-off transients.

One approach for facilitating successful nonintrusive monitoring is suggested by examining the spectral envelopes associated with the operation of a VSD. A typical power-electronic utility interface which might be used in a VSD is shown in Fig. 9. A three-phase, delta-connected utility service is rectified to provide a DC bus across a filter capacitor, C_f . The inductors in Fig. 9 model unavoidable utility line impedance. This DC bus powers a load modeled in Fig. 9 as the resistor R_L . In an actual variable-speed drive, the DC bus would provide power to an inverter that creates variable-frequency drive waveforms to vary the shaft speed of an induction or permanent-magnet synchronous machine.

A three-phase rectifier like the utility interface in Fig. 9 draws significantly distorted current waveforms i_a , i_b , and i_c on each of




Fig. 9 Power-electronic utility interface

the three utility lines. This distortion arises because of the switching action of the rectifier diodes, which can only turn on when a line-to-line voltage exceeds the bus-capacitor voltage. A typical experimental result is shown in the top trace of Fig. 10. This trace shows the line current, i_a , on one phase of a three-phase VSD for a 40-Hp induction machine driving a ventilating fan on the MIT campus. A delta-connected, 60-Hz, 208 line-to-line, RMS-VAC electrical service provides power for the VSD. For the data shown in the top trace of Fig. 10, spanning about two and one half line cycles of the voltage waveform, the drive was operating with approximately constant shaft speed. This nearly steady-state waveform contains significant harmonic content and can be approximated by a Fourier-series, i.e.,

$$\hat{i}_a = a_1 \sin(\omega_0 t) + a_5 \sin(5\omega_0 t) + a_7 \sin(7\omega_0 t),$$

where $\omega_0 = 2\pi 60$ radians-per-second, the base angular frequency for the utility, and the coefficients a_1 , a_5 , and a_7 indicate the amplitude of the fundamental, fifth, and seventh harmonic frequency components. The current \hat{i}_a is plotted in the bottom trace of Fig. 10 for the coefficients $a_1 = 13.03$, $a_5 = -8.59$, and $a_7 = 4.97$. Notice that, for these coefficients, the estimate \hat{i}_a is roughly equal to the actual measured current i_a shown in the top trace of Fig. 10. Other harmonics are certainly present in the measured current waveform, and the harmonic content may change slightly as the VSD responds to command changes. To a good approximation, however, as the VSD operates, drawing more or less input power, the input current approximately retains the shape shown in Fig. 10, but scales in magnitude.

This analysis leads to the expectation that a spectral envelope preprocessor observing the operation of a VSD should find time-varying content for the envelopes $a_1(t)$, $a_5(t)$, and $a_7(t)$, associated with the fundamental frequency and in-phase fifth and seventh harmonics, respectively. The three spectral envelopes plotted in Fig. 11 indicate that this is the case. These spectral envelopes are measured data acquired from the preprocessor during the turn-on transient of the 40-hp VSD. The amplitudes of the spectral envelopes are individually amplified by the data-acquisition system for convenient and accurate collection. Therefore, the plots in Fig. 11 correspond to the actual spectral estimates, scaled by a gain factor for each envelope. The drive begins with an "open-loop" spin-up to operating speed during the first 40 seconds of operation plotted in Fig. 11. From 100 seconds on, the drive is operating under closed-loop control as it attempts to regulate the pressure in a distant duct by varying fan speed.




Fig. 10 Measured and approximated VSD current




Fig. 11 Spectral envelope content of VSD

Spectral envelopes have great utility for nonintrusively monitoring the diagnostic condition and operation of variable-speed drives. The steady-state oscillations in nominal operation (after 100 seconds in Fig. 11) result from a poorly-tuned control loop in the HVAC system. These oscillations are relatively slow and easily missed by a casual inspection of the VSD control panel. They are very apparent, however, in the spectral envelopes. These oscillations could be detected with any of the event detectors described in the previous section, and could therefore be quickly detected by the NILM with transient event detection and spectral envelope preprocessing. For example, we have successfully detected power oscillations by measuring power variance about a moving-average mean.

Spectral envelopes could also conceivably be used to monitor the power consumption of continuously variable loads. The harmonic content in the current waveform tracks with the real power consumption of the VSD. That is, as the power consumption rises and falls, the seventh harmonic content, for example, also rises and falls. At target monitoring sites where a continuously variable load is the only load at the site that generates a particular harmonic, that harmonic could be used to track the operation of the load and disaggregate its power consumption and harmonic generation from the total signal at the service entry. If other loads are present that also generate overlapping harmonics, it might still be possible to disaggregate the individual contributions from each load or load type using orthogonal basis vectors associated with each load of interest; see [9] for some examples of this approach. This technique of using unique harmonic signatures to monitor variable-speed drives could be extended to any variable load that generates unique patterns in spectral envelopes.

VII Discussion

Nonintrusive electrical load monitoring is maturing into a reliable technology for inexpensively tracking the operation of electrical loads at a target site. Spectral envelope estimation and transient event detection and classification extend the capabilities of the NILM. Our field work indicates that these tools facilitate the monitoring of demanding sites with high rates of event generation and where steady-state information may be masked by power factor correction and load balancing.

Recent improvements in the performance and affordability of commercial DSP systems have made it possible to eliminate most of the relatively expensive custom hardware used in previous prototype nonintrusive monitors. This paper has introduced a new spectral envelope estimator implemented entirely in software on a DSP board. This estimator computes eight spectral envelopes in real time, and could be expanded entirely in software to compute nearly four times as many envelopes, or to perform other useful monitoring functions.

We have used the spectral envelope preprocessor to monitor the operation of a wide range of significant electrical loads in the laboratory and in the field. Transient event detectors described in this paper can be used to determine the operating schedule of loads, and also to monitor for certain kinds of faults or pathological conditions, e.g., the sustained steady-state oscillations seen in the 40-hp VSD. When a unique pattern of harmonic content can be associated with a particular load or load class, harmonic envelopes can be used in certain (previously intractable) cases to continuously and nonintrusively monitor the energy usage of variable loads.

Acknowledgments

The authors gratefully acknowledge the thoughtful comments of our reviewers and the valuable assistance of Deron Jackson, Jim Partan, and Dennis Evangelista. This research was funded in part by a National Science Foundation CAREER award and by MIT's Carl Richard Soderberg Career Development Chair. Essential hardware for this project was made available through generous donations from the Intel Corporation, Tektronix, and Hewlett-Packard.

References

- 1 Abler, C. B., "Spectral Envelope Estimation for Transient Event Detection," Master's thesis, MIT, 1998.
- 2 Basseville, M., and Benveniste, A., *Detection of Abrupt Changes in Signals and Dynamical Systems*, Springer-Verlag, 1980.
- 3 Hart, G. W., "Nonintrusive Appliance Load Monitoring," *Proceedings of the IEEE*, 1992.
- 4 Kahn, U. A., Leeb, S. B., and Lee, M. C., "A Multiprocessor for Transient Event Detection," *IEEE Transactions on Power Delivery*, 1997.
- 5 Kern, E. C., "Meter Sorts Load for Appliance Use," *Electrical World*, 1986.
- 6 Leeb, S. B., "A Conjoint Pattern Recognition Approach to Nonintrusive Load Monitoring," Ph.D. MIT, Department of Electrical Engineering and Computer Science, Feb. 1993.
- 7 Leeb, S. B., Kahn, U. A., and Shaw, S. R., Multiprocessing Transient Event Detector for Use in a Nonintrusive Electrical Load Monitoring System, Technical report, U.S. Patent Number 5,717,325, Issued Feb. 10, 1998.
- 8 Leeb, S. B., and Kirtley, J. L., A Transient Event Detector for Nonintrusive Load Monitoring, Technical report, U.S. Patent Number 5,483,153, Issued Jan. 9, 1996.
- 9 Leeb, S. B., Lesieutre, B. C., and Shaw, S. R., Determination of Load Composition Using Spectral Envelope Estimates, In 27th Annual North American Power Symposium.
- 10 Leeb, S. B., and Shaw, S. R., Harmonic Estimates for Transient Event Detection. In *Universities Power Engineering Conference*.
- 11 Leeb, S. B., Shaw, S. R., and Kirtley, Jr. J. L., "Transient Event Detection in Spectral Envelope Estimates for Nonintrusive Load Monitoring," *IEEE Transactions on Power Delivery*, 1995.
- 12 Siebert, William, *Circuits, Signals, and Systems*, The MIT Press, 1986.
- 13 EPRI test report, "Nonintrusive Appliance Load Monitoring System (nialms), Beta Test Results," EPRI number RS-108419, 1997.
- 14 Willisky, A., "A Survey of Design Methods for Failure Detection in Dynamic Systems," *Automatica*, 1976.