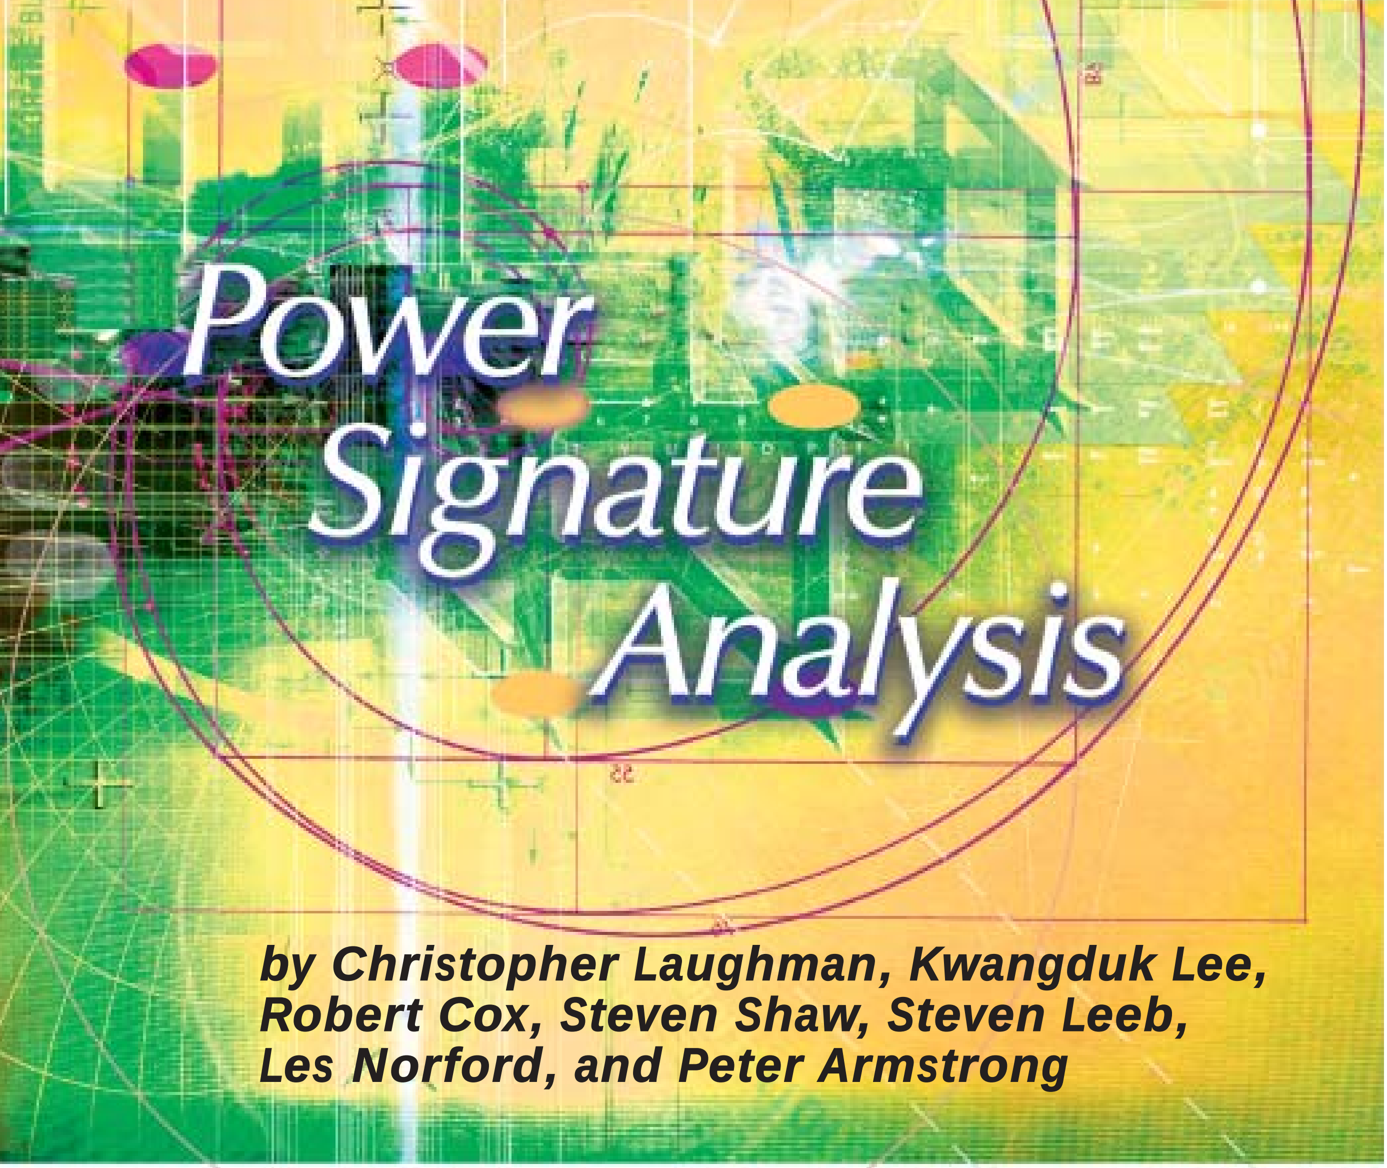

Power Signature Analysis

by Christopher Laughman, Kwangduk Lee, Robert Cox, Steven Shaw, Steven Leeb, Les Norford, and Peter Armstrong

THE BOOM IN COMMUNICATIONS RESEARCH AND DEVELOPMENT HAS MADE HIGHLY sophisticated means for gathering, moving, and exchanging information part of everyday life. Examples include the Internet, Bluetooth, and IEEE 802.11b networks, cellular, and satellite communications. However, the process of gathering and analyzing the data to send over this expanding web of information networks remains expensive for many applications. A building facilities manager may be able to configure a cellular phone to monitor the energy usage and operating schedule of an HVAC plant, but the usefulness of the information transmitted will generally be directly proportional to the complexity and size of the installed sensor array.

This difference between information gathering and information networking serves to reemphasize an established fact: though remote access to information and control inputs may be obtained easily and inexpensively via networking, access does not provide useful information without installation of a potentially expensive and intrusive sensor array. Mass production may ultimately reduce sensor cost, especially for solid-state or technologically advanced microelectromechanical system (MEMS) sensors. Installation expenses are likely to remain high, especially for temporary monitoring applications in which data is gathered for a brief window of time with a removable sensor network. In addition, the reliability of a monitoring system with many sensors may be reduced in comparison to a system with relatively fewer sensors.

1540-7977/03/$17.00©2003 IEEE

Nonintrusive load monitoring (NILM) can determine the operating schedule of electrical loads in a target system from measurements made at a centralized location, such as the electric utility service entry. In contrast to other systems, NILM reduces sensor costs by using relatively few sensors. NILM disaggregates and reports the operation of individual electrical loads like lights and motors using only measurements of the voltage and aggregate current at the utility service point of a building. It can also identify the operation of electromechanical devices in other kinds of power distribution networks. For example, NILM can determine the load schedule in an aircraft from measurements made only at the generator, or in an automobile from measurements made at the alternator/battery block. NILM can distinguish loads even when many are operating at one time.

NILM is an ideal platform for extracting useful information about any system that uses electromechanical devices. It has a low installation cost and high reliability because it uses a bare minimum of sensors. It is possible to use modern state and parameter estimation algorithms to verify remotely the "health" of electromechanical loads by using NILM to analyze measured waveforms associated with the operation of individual loads. NILM can also monitor the operation of the electrical distribution system itself, identifying situations where two or more otherwise healthy loads interfere with each other's operation through voltage waveform distortion or power quality problems.

Strategies for nonintrusive monitoring have developed over the last 20 years. Advances in computing technology make a new wealth of computational tools useful in practical, field-based NILM systems. This article reviews techniques for high-performance nonintrusive load and diagnostic monitoring and illustrates key points with results from field tests.

Background and Early Approaches

One of the earliest approaches to nonintrusive monitoring, developed in the 1980s at MIT by Fred Schweppe and George Hart, had its origins in load monitoring for residential buildings. Under Hart's ingenious scheme, the operating schedules of individual loads or groups of loads are determined by identifying times at which electrical power measurements change from one nearly constant (steady-state) value to another. These steady-state changes, known as events, nominally correspond to the load either turning on or turning off and are characterized by their magnitude and sign in real and reactive power. Recorded events with equal magnitudes and opposite signs are paired to establish the operating cycles and energy consumption of individual appliances.

This process of detecting steady-state changes provides the basis for a commercial version of this early work, as described in a 1999 article in *IEEE Computer Applications in Power.* The article describes a five-step process for load disaggregation through the detection of changes in aggregate power consumption.

- ✔ An edge detector is used to identify changes in steady-state levels.
- ✔ A cluster analysis algorithm is used to locate these changes in a two-dimensional signature space of real and reactive power (a ΔP-ΔQ plane). The signature space reduces the potentially complicated load transient data to a two-dimensional space of changes in power consumption with a pleasing and useful graphical interpretation.
- ✔ Positive and negative clusters of similar magnitude are paired or matched (especially for "two-state" loads that turn on, consume a relatively fixed power, and turn off).
- ✔ Unmatched clusters and events are paired or associated with existing or new clusters according to a best likelihood algorithm. This step is known as anomaly resolution.
- ✔ Pairs of clusters are associated with known load power consumption levels to determine the operating schedule of individual loads. This step uses information gathered during a training or survey phase in the building.

This five-step approach to nonintrusive load monitoring has the advantage of making intuitive sense and has been demonstrably successful in certain classes of buildings, such as residences. Recent field tests in a variety of other buildings have uncovered a number of limitations of the technique of examining steady-state changes in a signature space. Some of these limitations are well established, while others are relatively new.

> Field tests demonstrated the value of enhancements that enable an advanced nonintrusive load monitoring system to tackle complex monitoring environments.

The two-dimensional signature-space technique relies on at

> Savvy commercial and industrial facilities managers want near-real-time monitoring information in addition to batch results over an interval of several weeks.

least three key assumptions that limit its effectiveness. The first assumption is that different loads of interest exhibit unique signatures in the ΔP-ΔQ plane. They may not, especially in commercial and industrial facilities where the number and variety of loads is generally greater than in a residence. The two-dimensional signature-space becomes crowded with indistinguishable loads as the number and kind of loads increases. In a home, there may be only one 5 kW load, such as a hot water heater. In comparison, a light commercial facility with an HVAC plant, office equipment, and lighting may have several different loads that overlap ambiguously in the ΔP-ΔQ plane.

The second limiting assumption is that load composition is determined by steady-state power consumption. This requires waiting until transient behavior settles out so that steady-state values can be measured. We have found it difficult to find any suitable timescale in industrial or commercial environments that yields reliable steady-state measurements. A conservative estimate of steady-state power might require a long interval of near-constant demand. Requiring a long steady-state waiting time prevents the monitor from tracking rapid sequences of load activation. Some loads will be missed, and these may not be caught in the anomaly resolution phase. Short waiting times, on the other hand, may trigger measurements in the middle of load transients, resulting in spurious events in the cluster analysis phase. Field studies have shown that large HVAC loads such as fans and chillers might take from 30 s to several minutes to gradually spin up to their final operating speed. Other loads, such as variable speed drives (VSDs), may never settle to a steady-state. Furthermore, if the variations in power consumption are large enough, VSDs and similar loads could prevent the monitor from finding a steady-state consumption level or recording any events. Some industrial loads that ordinarily settle to steady-state conditions may fail to do so if they include poorly tuned controllers or other faults.

A third limitation is that most steady-state NILM systems process data in batch format using a day or more of stored data. This is not strictly implied by the five-step disaggregation procedure but is based on the assumption that near-real-time identification of load operation is not necessary. This limits the monitor to load survey and power scorekeeping

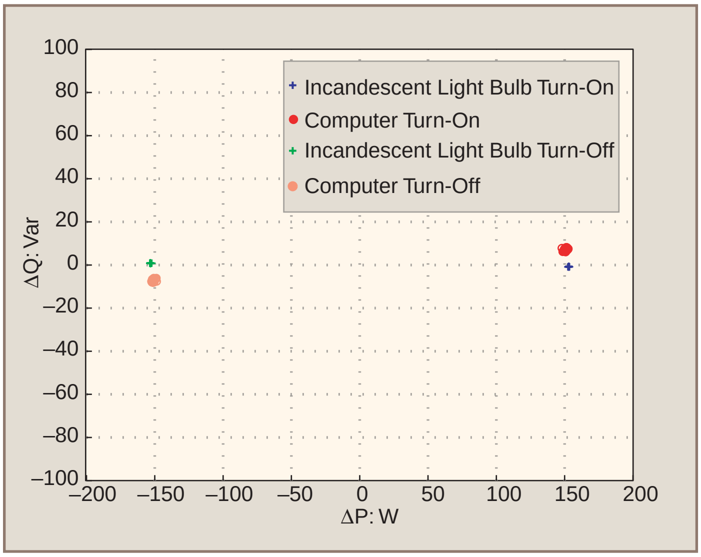

figure 1. Steady-state power consumption of a computer and a bank of incandescent lights; reactive power is plotted against real power.

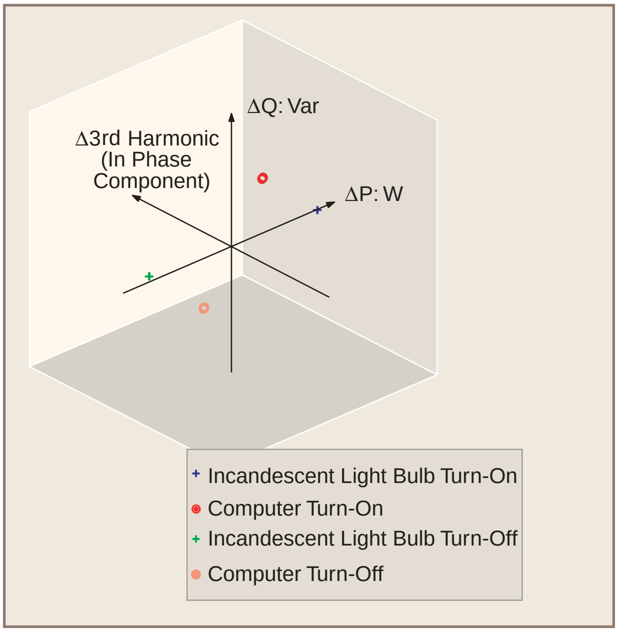

figure 2. Computer and incandescent light bulb turn-on and turn-off event clusters in the ΔP, Δ third harmonic, ΔQ coordinate system.

applications, excluding the vast potential for applications in real-time diagnostics.

Advanced Techniques for Monitoring

Steady-state monitoring techniques are successful in homes and small businesses because of the low event generation rate and number of loads at these sites. Medium to large size commercial and industrial facilities require a more sophisticated approach, due in part to high rates of event generation, load balancing, and power factor correction. Savvy commercial and industrial facilities managers want near-real-time monitoring information in addition to batch results over an interval of several weeks. Real-time identification is also essential if the load monitor is to serve as a "black-box" diagnostic monitor in transportation systems like aircraft or naval vessels. Field tests have demonstrated the value of a number of enhancements that enable an advanced noninstrsive load monitor to tackle complex monitoring environments.

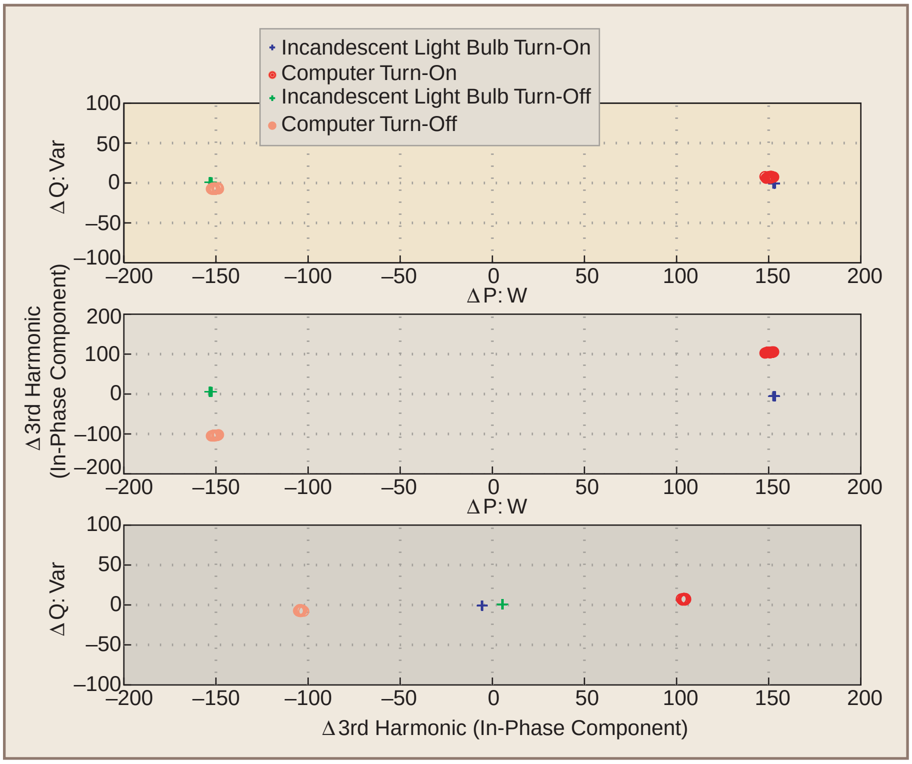

figure 3. Three cross-sections of the ΔP, Δ third harmonic, ΔQ coordinate system, illustrating the separation of clusters.

Higher Harmonics

Higher harmonics in the aggregate current signal can be used to distinguish loads with overlapping clusters in the ΔP-ΔQ signature space. Many loads draw distorted, nonsinusoidal currents due to their inherent physical characteristics or the presence of power electronics. Examples include office equipment (e.g., computers and copiers), actively controlled industrial equipment such as variable speed fans, and electroplating baths. Our advanced field monitoring system uses a phase-locked short-time Fourier transform of current waveforms collected at sample rates of 8,000 Hz or higher to compute spectral envelopes that summarize time-varying harmonic content. For a single-phase load, real and reactive power correspond to the envelopes of in-phase and quadrature current drawn by the load relative to the voltage. The short-time Fourier transform computes estimates of the real, reactive, and higher frequency components of the current.

Figure 1 shows a collection of turn-on and turn-off events recorded at a site with personal computers and a bank of incandescent lamps. The loads are practically indistinguishable in the ΔP-ΔQ signature space because they consume essentially the same real and reactive power. However, typical computer power supplies draw signature third harmonic currents. The loads are easily separable by examining higher harmonics.

Figures 2 and 3 illustrate the relative ease of distinguishing individual loads in a three-dimensional space with axes denoting changes in real power, reactive power, and third harmonic. Our advanced load monitor routinely examines harmonic content up to and including seventh harmonic and can be customized to examine higher harmonics as necessary.

Transient Detection

Our advanced load monitor recognizes individual loads based on distinctive load transient shapes. The transient behavior of a typical load is intimately related to the physical task that the load performs. For example, the turn-on transients associated with a personal computer and with an incandescent lamp are distinct because charging capacitors in the computer power supply is fundamentally different from heating a lamp filament. Overall transient profiles tend to be preserved even in loads that use active waveshaping or power factor correction. Most loads observed in the field have repeatable transient profiles, or at least sections of the transient profile that are repeatable. Transient-based recognition permits near-real-time identification of load operation, especially turn-on events.

Transients are identified by matching events in the incoming aggregate power stream to previously defined transient signatures, or exemplars. Exemplars can be determined, for example, by a one-time direct observation of the device in question or by previous training in the laboratory. Pretraining has proven to be a reasonable approach for very repeatable loads that show up in large quantities, such as fluorescent lamp ballasts. The exemplar may be composed of multiple parts for transients with a number of distinct sections. Each section of the exemplar can be shifted in time and offset to match incoming transient data. In addition, an overall gain may be applied to all sections of the exemplar to achieve a better fit. Each event detected is compared to the full set of exemplars by using a least squares criterion to select the appropriate shifts and gains. The match with the lowest residual norm per number of points is then compared to a threshold. If the fit is good enough, the event is classified as a match to the exemplar. If not, the event is left unclassified. Correct classification of overlapping transients is possible using properly designed exemplars.

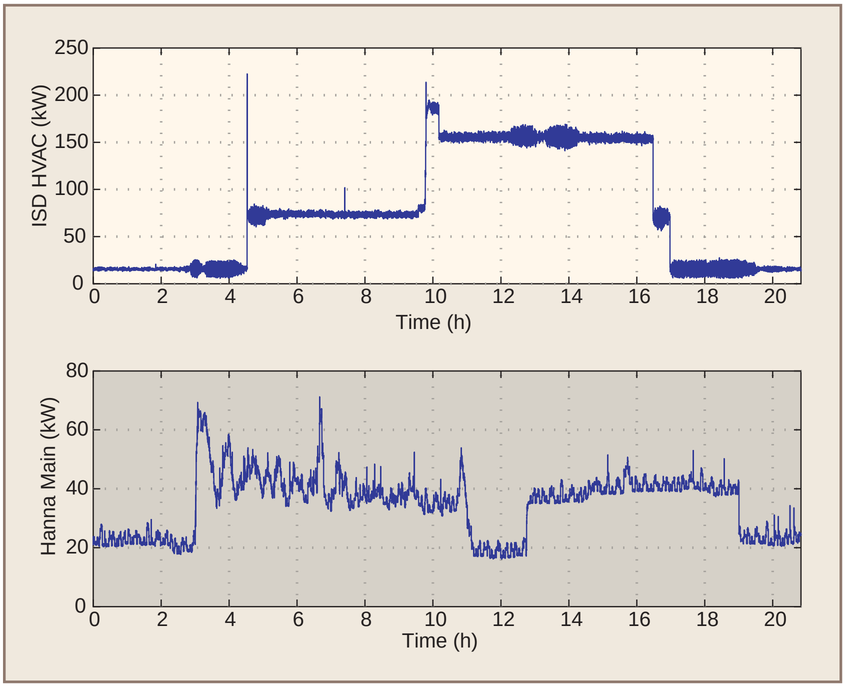

figure 4. Plots of the aggregate energy consumption collected at the HVAC service entry in a California county government building and at the main service entrance of a California elementary school.

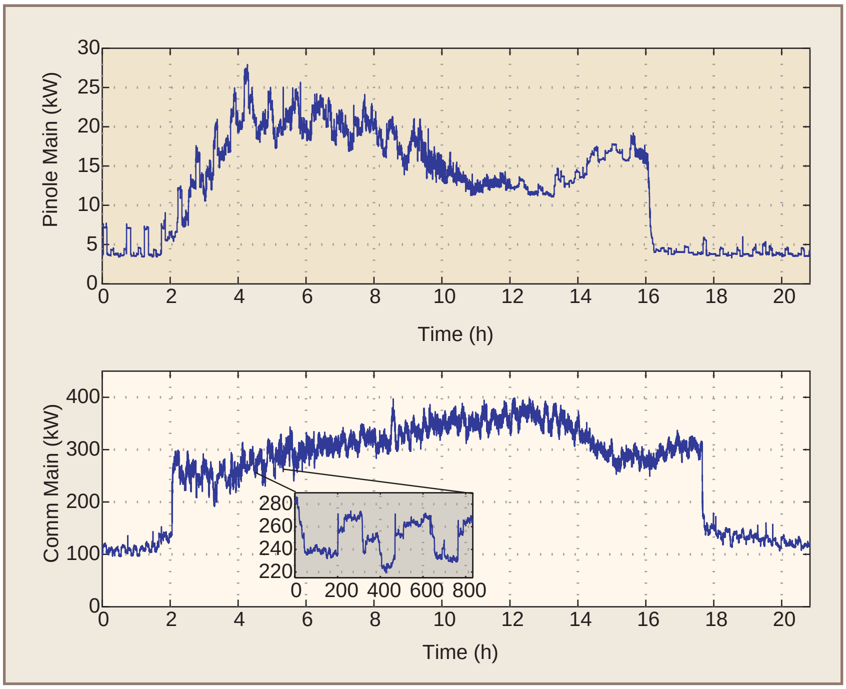

figure 5. Plots of the aggregate energy consumption collected at the main service entries of a California middle school and a communications service building. The inset plot illustrates the high density of events; this data in this plot is averaged in order to enhance the visibility of the individual events.

The value of examining transient information can be understood graphically by examining Figures 4, 5, and 6. Figures 4 and 5 show the spectral envelopes corresponding to the fundamental, in-phase component of current (essentially, real power) demanded in four different buildings in California. The top graph in Figure 4 shows the power consumed solely by the HVAC panel in a building (ISD) in Los Angeles County. The bottom trace of Figure 4 and the top trace of Figure 5 show power consumption in two different public schools near San Francisco (Hanna and Pinole, respectively). Finally, the bottom trace in Figure 5 shows the total power consumption in another building in Los Angeles County (Comm); the inset in this plot is a magnified portion of the same waveform, illustrating the high density of events. These four traces show a progression in increasing event generation in a building HVAC panel using steady-state signature detection. As the complexity and variation on the electrical grid increases, it becomes more difficult to tune a steady-state change detector to function at all. For buildings with behavior nearing that of the Comm building, it becomes essential either to augment the steady-state approach or to discard it completely in favor of transient identification. By enabling the monitor to deal with higher rates of event generation, a single nonintrusive load monitor using transient event detection can generally monitor a comparatively more complex building power network than with steady-state change detection alone.

In addition to better load disaggregation, transient analysis can provide diagnostic information. Diagnostics can be developed by exploring the

Nonintrusive load monitoring can determine the operating schedule of electrical loads in a target system from measurements made at a centralized location, such as the electric utility service entry.

relationship between the electrical transient and the physics of the load. Specifically, a load model can often be developed that relates the shape of the electrical transient to the physical parameters of the load. The load model is often a differential equation, although we have experimented with other choices. It is often possible to deduce the physical condition of the load by examining load model parameters when the transient data is rich enough to identify the load model.

Figure 6 illustrates the diagnostic capabilities of our advanced load monitor. The numbers displayed in the text window below the graphs in Figure 6 show estimated parameters of an induction machine model. The solid lines in the graph represent raw data, while the points show the results of a simulation with the estimated parameters. The close agreement shown in the figure indicates that the parameters and model accurately predict the response of the actual load. In addition to developing these physically based models, we have also developed techniques that can accurately estimate model parameters without a good initial guess. These estimation methods enable a building manager to trend load parameters and use parameter information to predict impending faults. In the case of an induction motor, such faults might include shorted motor windings, broken rotor bars, and especially failures in a mechanical system attached to the shaft, e.g., a slipping belt.

Disaggregating Continuously Variable Loads

A third benefit obtained through transient identification and spectral analysis is that our advanced load monitor has the ability to monitor buildings with continuously variable loads. Many loads, such as the VSDs used in HVAC fans, servo their electrical power consumption under the influence of an active closed-loop controller. Transient detection can simplify the tracking of these and other continuously variable loads in several ways.

Figure 7 shows three spectral envelopes associated with the operation of a 40 hp VSD in an HVAC plant on the MIT campus. The top trace in this figure shows the real power demanded by a variable-speed fan drive in an HVAC system. The drive begins with an open-loop spin-up to operating speed during the first 40 s of operation. This open-loop spin-up is repeatable because a microprocessor controls the startup profile every time the drive is activated. NILM can recognize that a VSD is active in the building using transient recognition, but transient recognition does not provide a means for continuously tracking the power consumption of a variable load like a VSD.

When NILM recognizes a continuously variable load like a VSD, we have found that it is often possible to disaggregate the variable load by carefully examining the spectral envelopes. The VSD connects to the utility through a three-phase, delta-connected rectifier. The three-phase rectifier draws distorted, pulsatile current waveforms from each of the three utility phases. As shown in the lower two traces in Figure 7, this rectifier set has the effect of causing the VSD to create characteristic traces not only in real power but also in the fifth and seventh harmonics, as illustrated in the middle and bottom traces respectively. These higher harmonic traces have

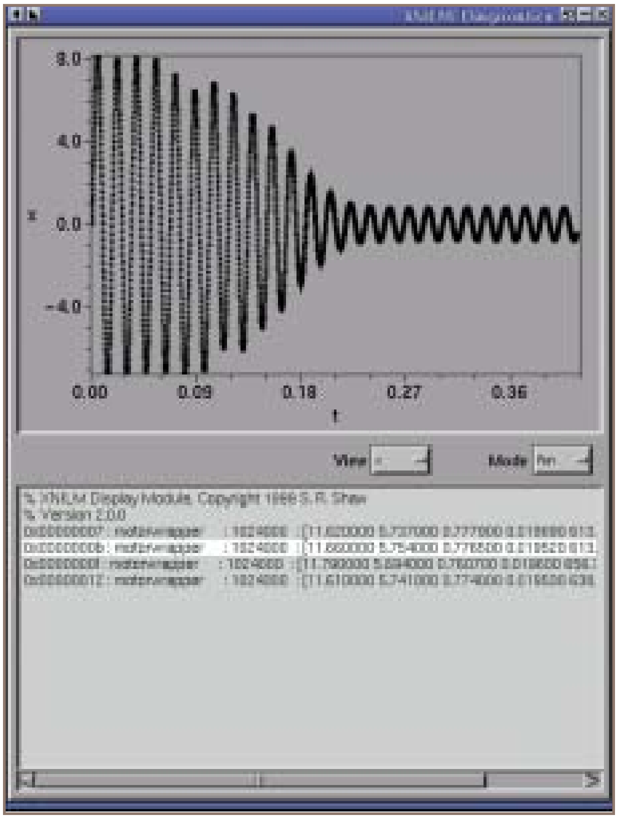

figure 6. Startup transient for an induction machine.

envelopes that roughly track the shape of the real power trace. If the VSD is the only load at a target site that generates the bulk of a particular harmonic signal (such as the seventh harmonic) in the aggregate data, this particular harmonic signal can be functionally mapped and subtracted from the real power spectral envelope in order to remove the continuously varying component of the real power signal due to the VSD from the aggregate power signal. This makes it easier to identify the remaining active loads in the aggregate power signal. In general, the higher harmonic information can be used both to track the energy consumption of the variable load and also to disaggregate this variable load from the power consumption of the remaining loads. There are many possible variations on this technique of refining the mappings between a particular harmonic and other harmonics; these techniques allow one group of loads to be disaggregated from another.

Additionally, harmonic signatures can provide direct diagnostic information even without parameter estimation. Returning to Figure 7, the drive is operating under closed loop control as it attempts to regulate the pressure in a distant duct by varying fan speed starting at 100 s. The control loop is poorly tuned, and the drive exhibits a slow, large amplitude oscillation in steady state. This pathology is not easily caught in a brief inspection of the drive, as the oscillation is fairly slow; nevertheless, the identification of this failure mode is important because it wastes energy and wears the mechanical components of the drive. This condition is easily caught by an advanced nonintrusive load monitor that examines transients and spectral envelopes.

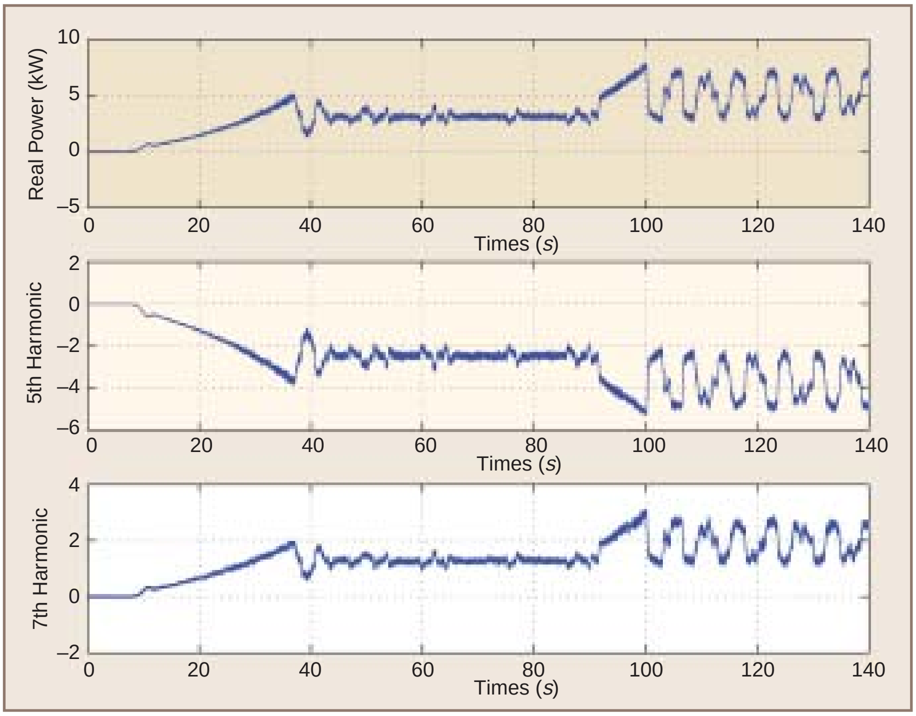

figure 7. Startup transient of a variable speed drive.

Applications and the Future

We have implemented all of these analytical enhancements for nonintrusive monitoring on an inexpensive computer platform for use in the field. A suite of NILM tools, including sophisticated model-based diagnostic algorithms that track (and could in principle trend) model parameters to determine the health of critical loads, has been developed under the Linux operating system environment. All of these software tools run on a Pentium-class personal computer (300 MHz clock or higher) with a PCI-bus data-acquisition card; any relatively inexpensive personal computer, laptop, or embedded system like a PC104 chassis could be used to develop a modern monitor that performs transient event detection. We have used all of these platforms with success in buildings and transportation systems.

Current experiments in the field are aimed at enhancing the diagnostic capabilities of the monitor, integrating it into building energy management and control systems, and improving it as a load monitor for conducting surveys and energy-usage scorekeeping. Even in very modern buildings with sophisticated control and consumption metering systems, we have found that the monitor is invaluable to sophisticated facilities managers. As the monitor requires virtually no additional wiring effort and few installed sensors, it provides invaluable building operating history and health information in a package that is both easy to install and highly reliable.

In the future, we expect to introduce new event detection and classification schemes in which load disaggregation and usage tracking are increasingly accurate. The best approaches appear to blend all the tools described above to some degree. For example, the highest accuracy in usage disaggregation may be obtained by combining the steady-state and transient approaches to event recognition. Additionally, the steady-state signature scheme can obviously be enhanced by working in a larger signature space with harmonic data. Yet another combination of steady-state and transient identification methods may be used to alleviate the problem of defining steady-state intervals by using transient event detection to tease apart overlapping transients which occur at high rates. Furthermore, the transient detection approach may be minimally successful when transient shapes are masked or distorted by background, periodic oscillations in the building; by examining changes in the steady-state sections of the aggregate data stream, the transient event detector may be more successful in looking for most likely shapes in a cluster of overlapping or distorted transients. We are currently developing arbiters that use many different approaches to identify events with the highest possible accuracy and precision.

Because the advanced load monitor uses a conventional computer platform, it is easily connected with communications networks, including the web. An example from a monitor installed on the MIT campus is shown in Figure 8. The Netscape window shows a turn-on transient associated with a washing machine on top of an 11 kW base load; this can be viewed via a remote connection to a computer monitoring a campus dormitory containing a laundromat. The cusp-like steady-state, following the initial acceleration of the agitator motor to operating speed, is due to the steady-state oscillation of the agitator. Other load transients recognized by the monitor, including several dryers, are listed in the text window below the graphs. The items in this list are hypertext links corresponding to recorded events; these links can be selected in order to view a record for each indicated transient, including time, energy consumption, and diagnostic parameters (e.g. heater resistance, motor parameters like stator resistance, magnetizing and leakage inductances, etc.). We anticipate that the advanced NILM technology will provide an ideal platform for remote monitoring and for customizing network-accessible information packages for utility customers and building occupants.

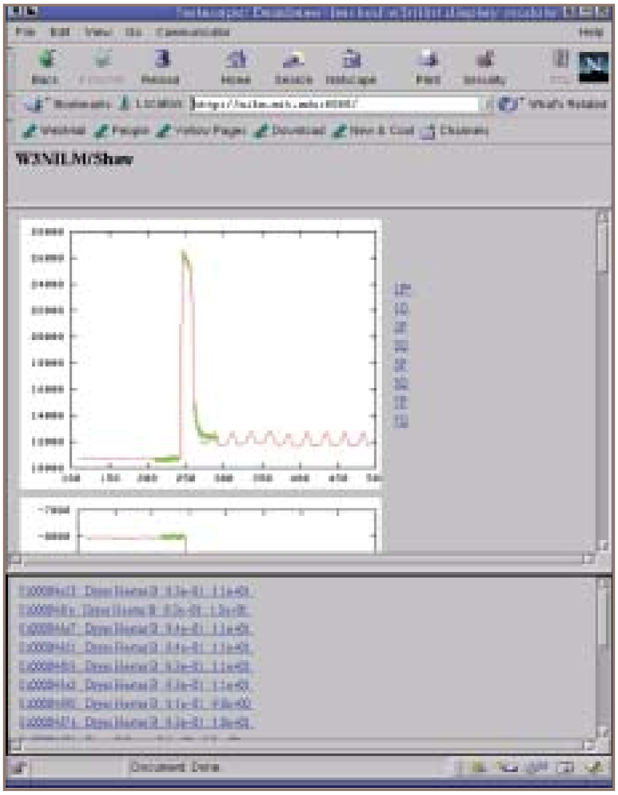

figure 8. Noninstrusive monitoring on the Web.

Acknowledgments

This research was funded by the California Energy Commission, a grant from the Grainger Foundation, and the Ford Motor Company.

Further Reading

"Wired and wireless," *Technol. Rev. Mag.*, vol. 104, no. 5, pp. 40-82, Jun. 2001.

S. Drenker and A. Kader, "Nonintrusive monitoring of electrical loads," *IEEE Comput. Appl. Power*, vol. 12, no. 4, pp. 47-51, 1999.

G.W Hart, "Nonintrusive appliance load monitoring," *Proc. IEEE*, vol. 80, no. 12, pp. 1870-1891, Dec. 1992.

S.B. Leeb, S.R. Shaw, and J.L. Kirtley, Jr., "Transient event detection in spectral envelope estimates for nonintrusive load monitoring," *IEEE Trans. Power Delivery*, vol. 10, no. 3, pp. 1200-1210, Jul. 1995.

L.K. Norford and S.B. Leeb, "Nonintrusive electrical load monitoring in commercial buildings based on steady state and transient load-detection algorithms," *Energy and Buildings*, vol. 24, pp. 51-64, 1996.

S.R. Shaw, C.B. Abler, R.F. Lepard, D. Luo, S.B. Leeb, and L.K. Norford, "Instrumentation for high performance nonintrusive electrical load monitoring," *ASME J. Solar Energy Eng.*, vol. 120, no. 3, pp. 224-229, Aug. 1998.

S.R. Shaw and S.B. Leeb, "Identification of induction motor parameters from transient stator current measurements," *IEEE Trans Ind. Electron.*, vol. 46, no. 1, pp. 139-149, Feb. 1999.

S.R. Shaw, R.F. Lepard, S.B. Leeb, and C.R. Laughman, "A power quality prediction system," *IEEE Trans. Ind. Electron.*, vol. 47, no. 3, pp. 511-517, Jun. 2000.

Biographies

Christopher R. Laughman received S.B. and M.Eng. degrees from the Massachusetts Institute of Technology (MIT). He is a research engineer in the Laboratory for Electromagnetic and Electronic Systems. His current research includes the investigation of modeling and parameter estimation problems in buildings and transportation systems.

Steven R. Shaw received his doctoral degree from MIT in electrical engineering and is currently an assistant professor in the department of Electrical and Computer Engineering at Montana State University in Bozeman. He is interested in sensors, instrumentation, modeling, numerical and computational methods associated with control, and measurement problems.

Steven B. Leeb received S.B., S.M., E.E., and Ph.D. degrees from MIT. He has been a member of the MIT faculty in the Department of Electrical Engineering and Computer Science since 1993. He serves as an associate professor in the Laboratory for Electromagnetic and Electronic Systems. He 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 an IEEE Senior Member.

Leslie K. Norford earned his Ph.D. in mechanical and aerospace engineering at Princeton University. He is an associate professor of building technology in the Department of Architecture at MIT, where he has been a member of the faculty since 1988. His research interests concern monitoring the performance of mechanical and electrical equipment in buildings, optimization techniques as applied to the design of buildings and their mechanical systems, simulation of the performance of building equipment, and efforts in support of sustainable buildings in developing countries.

Robert W. Cox, Kwangduk D. Lee, and Peter Armstrong are graduate students pursuing their doctorates in the departments of Electrical Engineering, Mechanical Engineering, and Building Technology, respectively.