AC-02-12-4

Monitoring HVAC Equipment Electrical Loads from a Centralized Location—Methods and Field Test Results

Dong Luo, Ph.D.
Student Member ASHRAE

Leslie K. Norford, Ph.D.
Member ASHRAE

Steven R. Shaw, Ph.D.

Steven B. Leeb, Ph.D.

ABSTRACT

This paper reports recent work to determine useful information about component-level HVAC electrical loads—not from submeters, which are accurate, rarely installed, and relatively expensive, but instead from one or more centralized locations in a building's electrical distribution system. The work includes laboratory tests with real-building data and field tests made with low-cost hardware capable of the rapid sampling needed for load disaggregation. Results indicate that building electrical signals are often quite complex, that individual loads can indeed be detected with reasonable reliability, that more work is required to automate the process of tuning the detection algorithm, and that there are benefits to analyzing turn-on/turn-off events at multiple sampling rates to minimize trade-offs between detection sensitivity and false alarms.

INTRODUCTION

Accurate and affordable information about HVAC electrical loads is of value to many individuals and organizations involved in providing HVAC services: facility managers, who would like to minimize operating costs and the costs and down-time associated with repairs; electric utilities and service providers, who need accurate load models to most economically generate, transmit, and distribute power; and energy service companies and building owners, who would like inexpensive means to verify savings from energy-efficient improvements. Electrical-power information can also be used for power-quality monitoring and for analyzing loads other than HVAC, including lights and process equipment.

The demands for both reasonable accuracy, suitable for the task at hand, and moderate cost, in keeping with achievable benefits, are not easily met. The required accuracy varies with specific monitoring goals, graduated as follows:

1. detection of on-off switching events, to determine whether equipment is operating in accordance with the expected schedule and whether it responds to a control signal;
2. measurement of power magnitudes at the time of on-off switches and throughout the operating cycle, to compute energy consumption; and
3. measurement of changes in power or changes in the frequency of operation of equipment, both relative to normal operation, as a basis of detecting faulty operation.

The first level can be easily achieved with current transducers, now in limited use in buildings to effectively provide an echo of a control signal and thereby verify that a fan or pump has turned on or off on command. This simple technology is not universally applied, suggesting some ambivalence about cost versus benefit. Going to the second and third levels requires more expensive power metering rather than binary (on-off) current information. The third level, fault detection and diagnosis, also requires development of methods to analyze changes in electrical-power data and to relate these changes to normal operating patterns.

This paper describes recent work on an electrical-load monitor capable of obtaining electrical-power data at a cost lower than submeters. This monitor has its origins in residential buildings, where it was designed to be installed in lieu of the standard revenue meter and hence did not cross the utility-customer boundary. In that sense it was not invasive and was therefore known as a non-intrusive load monitor, or NILM. The analysis method for the residential meter, detailed in Hart (1992) and now in commercial production, is based on

Dong Luo is a senior engineer with United Technologies Corporation. **Leslie K. Norford** and **Steven B. Leeb** are associate professors at the Massachusetts Institute of Technology, Cambridge, Mass. **Steven R. Shaw** is an assistant professor at Montana State University, Bozeman.

pinpointing the times at which a near-constant series of electrical power measurements changes to another near-constant series. Changes are characterized by their magnitude in real and reactive power. Changes of near-equal magnitude and opposite sign are paired to establish the operating cycles and energy consumption of individual residential appliances.

This approach faces several limitations in commercial buildings and, increasingly, in houses:

1. electrical noise generated by power electronics, which makes it difficult to establish steady-state conditions and, once steady state is defined, also generates changes in power that must be analyzed as events even though they are not of interest;
2. overlapping on-off events that may mask individual changes; and
3. time-varying electrical power demand by individual components.

The first two limitations have been overcome in two ways. First, as described in detail elsewhere (Leeb 1993; Leeb et al. 1995; Norford and Leeb 1996; Leeb and Kirtley 1996; Leeb et al. 1998), a method has been developed to measure and analyze the short-term dynamic electrical pattern associated with the start-up of a piece of equipment to aid in identifying this equipment in an environment that is either noisy or rich with information. This approach computes spectral envelopes at the fundamental and higher harmonic frequencies, detects rapid changes in these envelopes, and compares the patterns of such changes with libraries of known patterns, generated for either classes of equipment (induction motors, rapid-start lamps) or individual components (a grinder, for example). Another approach to analyzing start-up transients as a means of non-intrusive load detection is presented in Deschizeau et al. (2000).

Detection and analysis of start-up transients hold the promise of a powerful approach to fault detection, requiring only short-term, focused, and robust power analysis rather than more extended computation of changes in power consumption under known loading conditions. Recent efforts in this area rely on submetered rather than centralized power measurements and are described in Shaw et al. (2002).

Second, a less powerful, more easily implemented, and ultimately complementary effort detects changes in power levels in a manner that is similar in spirit to the original residential concept but more effective in noisy electrical environments encountered in commercial buildings. Norford and Leeb (1996) showed examples of how one aspect of this work, median filtering, can improve signal quality. Abler et al. (1998) introduced the signal processing algorithms. Work reported in this paper builds on the initial efforts of Hill (1995), provides a full discussion of material introduced in Abler et al. (1998), and presents a significant extension to cope with signals of interest that are small relative to noise. This approach has immediate application in a number of areas, including fault detection and diagnosis, and also provides a means of triggering the more computationally intensive transient-event detector.

This paper first describes an appropriate technique for detecting turn-on/turn-off events in noisy electrical power signals. The basic method precedes a number of refinements. The refined approach is then melded with an oscillation detector, to shield the on-off detector from the impact of oscillatory power signals and to analyze those oscillations as indicators of poorly tuned HVAC controllers. Next, the on-off detector is extended to operate at multiple sampling rates, found necessary in order to reliably discern cycling patterns for a two-stage reciprocating chiller in a whole-building electrical signal. Finally, the paper compares centralized power measurements with submeters, briefly describes the monitoring hardware, and offers conclusions. Most of the data used in this work were collected at the test building used for an ASHRAE-sponsored research project on fault detection in HVAC systems, RP 1020; the test building is described more fully in Norford et al. (2000, 2002).

BASIC DETECTION ALGORITHM—GENERALIZED LIKELIHOOD RATIO (GLR)

Figure 1 shows representative electrical-power data collected at the ASHRAE test building. Electrically noisy commercial buildings have led to a large number of false alarms when detecting turn-on/turn-off events via changes in steady-state power (Norford and Mabey 1992). A statistical algorithm more reliable and powerful than a simple trigger based on deviation from the mean has been developed by extending the generalized likelihood ratio (GLR) (Basseville and Nikifirov 1993).

The GLR detection algorithm calculates a decision statistic from the natural log of a ratio of probability distributions before and after a potential change in mean:

$$S_j^k(\mu_1) = \sum_{i=j}^{k} \ln \frac{P_{\mu_1}(y_i)}{P_{\mu_0}(y_i)}, \tag{1}$$

where

y_i = sampled variable at time i;

μ_0, μ_1 = mean values of the sampled sequence before and after the event, respectively;

$P_\mu(y_i)$ = probability density function of the sampled sequence $y_i(i = j, \ldots, k)$ about the mean value μ ;

$S_j^k(\mu_1)$ = detection statistic, which is the log likelihood ratio of the joint frequency function for the independent variables y_i during time j to k about μ_1 and μ_0 .

The behavior of the detection statistic in the absence or presence of a step change in electrical power can be easily understood. Suppose that the function is used to test whether a chiller of known electrical power has turned on. Before the event to be tested, power data, collected at a measurement point that includes electrical service to the chiller, are distributed about the pre-event mean power level. If the chiller

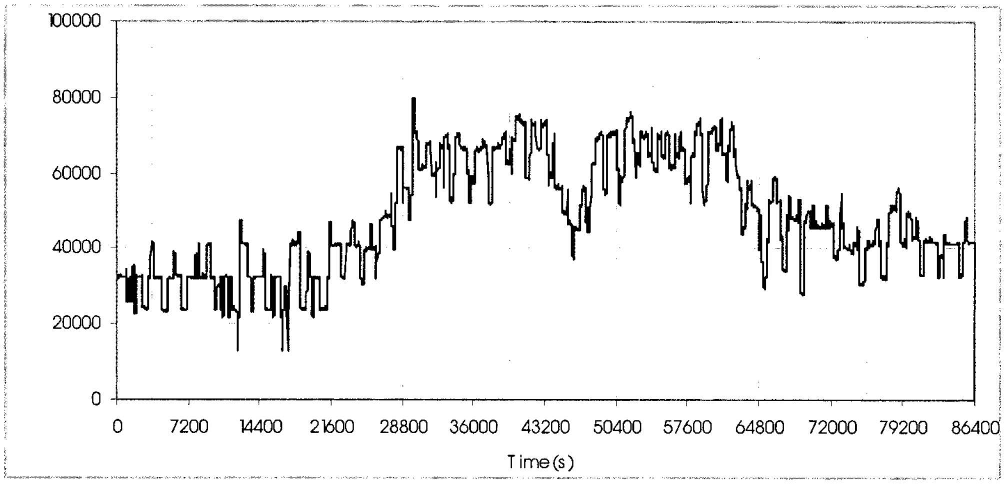

Figure 1 *Whole-building electrical power, sampled at 24 Hz. Data were collected at the electrical-service entry to the test building used in ASHRAE RP-1020.*

does not turn on, the power measurements will not change. The probability density function (PDF) in the numerator of Equation 1, which is centered on a post-event mean that includes the chiller power, will be very small, while the PDF in the denominator will be much larger, because data points will still be clustered about the pre-event mean. In other words, there is a very small probability that the measured power level would be associated with the operating chiller. The probability ratio is therefore very small and the log of the ratio is large and negative. As the magnitude of the hypothesized change in power decreases, the ratio approaches a value of 1.0 in the absence of a power change and the natural log approaches a value of 0. By contrast, if the chiller turns on, the PDF in the numerator will be large, because data points will be centered about the hypothesized mean, and the PDF in the denominator will be small. The ratio, therefore, is large.

There are two independent variables, the change time and the mean value after the change, which leads to a double maximization of the detection statistic S_j^k:

$$g_k = \max_{l \le j \le k} \ln \hat{\Lambda}_j^k = \max_{l \le j \le k} \sup_{\mu_1} S_j^k(\mu_1) , \tag{2}$$

where

$\hat{\Lambda}_j^k$ = estimate of the upper bound of the ratio of the joint frequency function about the post-event mean value μ_1 for a given pre-event mean μ_0, within the window [j, k].

sup = supremum, i.e., the least upper bound of S_j^k over [j, k] about the mean value μ_1 with reference to the known mean μ_0 before the change.

In other words, the event identified by the maximum probability ratio is found by searching for the time j and the corresponding average μ_1 in each subwindow j-k in the current detection window 1-k.

In some cases, the minimum magnitude of the change in power for a given component is known in advance. The minimum value of the change of parameter μ, designated as V_m, can be used in the search of μ_1, i.e., $|\mu_1 - \mu_0| \ge V_m$. Equation 2 can then be rewritten as:

$$g_k = \max_{l \le j \le k} \sup_{|\mu_1 - \mu_0| \ge V_m \ge 0} S_j^k(\mu_1) . \tag{3}$$

Because noise in electrical power of equipment follows the normal distribution, according to the central limit theorem (Rice 1988), noise in the total power data of a system consisting of components independent in power consumption can also be described by the normal distribution. Therefore, the deviation of the sampled total power data y_i from the calculated mean μ can be represented by the normal distribution, $N \sim (0, \sigma^2)$. Here, σ stands for the standard deviation of the total power data of the monitored system.

For an independent Gaussian sequence, the probability density function is

$$P_\mu = \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{-(y_i - \mu)^2}{2\sigma^2}} . \tag{4}$$

The detection statistic can then be derived as

$$S_j^k = \frac{\mu_1 - \mu_0}{\sigma^2} \sum_{i=j}^{k} \left(y_i - \frac{\mu_1 + \mu_0}{2} \right) . \tag{5}$$

Let $V = \mu_1 - \mu_0$, the change in mean power signal after an on-off event. Then

$$g_k = \max_{l \le j \le k} \sup_{|V| \ge V_m > 0} \sum_{i=j}^{k} \left[\frac{V^*(y_1 - \mu_0)}{\sigma^2} - \frac{V^2}{2^*\sigma^2} \right] \quad (6)$$

and

$$\hat{V}_j = \left(\frac{1}{k-j+1} \sum_{i=j}^{k} |y_i - \mu_0| - V_m \right)^+ + V_m \quad (7)$$

where $\hat{V}_j$ is the value of V at which g_k reaches its maximum,

$$\hat{V}_j = \begin{cases} V_m & \text{if} \quad \frac{1}{k-j+1} \sum_{i=j}^{k} |y_i - \mu_0| < V_m \\ \frac{1}{k-j+1} \sum_{i=j}^{k} |y_i - \mu_0| & \text{if} \quad \frac{1}{k-j+1} \sum_{i=j}^{k} |y_i - \mu_0| \ge V_m . \end{cases}$$

Then g_k can be rewritten as

$$g_k = \max_{l \le j \le k} \sum_{i=j}^{k} \left[\frac{\hat{V}_j^*(y_1 - \mu_0)}{\sigma^2} - \frac{\hat{V}_j^2}{2^*\sigma^2} \right]. \quad (8)$$

If $V_m = 0$, meaning that a change of mean of any magnitude is of interest or no information about the minimum expected change is available in advance, then

$$g_k = \frac{1}{2^*\sigma^2} \max_{l \le j \le k} \frac{1}{k-j+1} \left[\sum_{i=j}^{k} (y_i - \mu_0) \right]^2 . \quad (9)$$

The magnitude of g_k increases with the change in power and the abruptness of the change. A value above a threshold indicates an on-off event of potential interest.

Training the Parameters of the GLR Algorithm

The GLR detector requires that four parameters be trained for a given application:

1. the length of the pre-event averaging window;
2. the length of the detection window;
3. the threshold for the detection statistic; and
4. the standard deviation (or variance) of the power data.

Guidelines follow for selecting the parameters in one or more of three ways: *a priori* experience, tuning the parameters on the basis of one-time measurements, or automatically adjusting the parameters. The first two guidelines concern the size of sliding windows, as expressed in a number of samples. These guidelines have been shown to be reasonable for a range of sampling rates. Guidelines for the sampling rate itself will be presented later in this paper.

The Length of the Pre-event Averaging Window. Because multiple power changes occur in sequence in HVAC systems, the mean before the change must be continuously updated with each new data point. The length of the data window used to estimate the pre-event average power has a profound impact on the detection statistic. The GLR algorithm with a short window will yield a large detection statistic when the incoming data include spikes, which will be a source of false alarms if the detection statistic exceeds its threshold value. Further, the GLR algorithm with a short window may also miss events that occur gradually over several data points. On the other hand, a GLR detector with a long window will not find multiple abrupt changes that are close to each other in time.

On the basis of trial-and-error tuning, the appropriate length of the pre-event averaging window was 10 data points for a set of one-minute-averaged data that recorded electrical power for four air handlers, each consisting of a supply and a return fan. At the ASHRAE test building, this window was set to be six data points, sampled at one-second intervals. An estimate of a suitable window length can be made on the basis of the characteristics of the HVAC components, how often they are switched on or off, the status of the electrical facility, and the typical power profile of the system. Observations made in this research indicate that the upper limit for the length of the pre-event averaging window should be no longer than the interval between two major consecutive events. This interval can be estimated from a basic knowledge of the monitored HVAC system. Major individual components usually have minimum on and off times, to protect against deterioration due to frequent on/off switching. Moreover, systems of components are or can be controlled to operate in a sequence with specified time intervals, as in turning on or off air-handler fans. Note that this requirement is based in the time domain and is converted to a number of samples on the basis of the sampling rate.

As a lower limit, the pre-window should never be shorter than significant electrical noise spikes and ideally should be longer than the duration of the start-up period of each single component in the system. However, this condition cannot always be met for a real system, because the start-up of a VSD fan can be as long as 15 minutes while the interval between the on/off transitions of two components can be shorter than this. In such cases, other approaches are needed, such as the multi-rate sampling technique discussed later. Practical experience has shown that the window should contain at least four data points.

Without violation of the above basic rules, the pre-window should be kept short. If there is a range of possible intervals between the lower and upper bound, a value at or near the lower bound should be selected. For a given system, the window length for detection seems to be consistent across different seasons and equipment operating conditions, as demonstrated by extensive tests to detect HVAC equipment on-off events in the ASHRAE 1020-RP test building (Shaw et al. 2002; Norford et al. 2000, 2002).

If the multi-sampling rate detection algorithm (to be described later) is used, the number of data points in the window should be determined from the above time-domain guidelines and the fastest sampling rate. The same number of data points are used at all sampling rates in the multi-rate algorithm.

The Length of the Detection Window. The above basic limits for the pre-event averaging window also apply for the post-event detection window—i.e., it should not be longer than the interval between two consecutive events, never shorter than a disturbance, and ideally not shorter than the duration of a start-up transient process. On the other hand, unlike the pre-event averaging window, which is used to achieve a stable mean as the reference for coming events, the post-event detection window is intended to be sensitive to events yet robust to disturbances. A shorter post-event detection window is more sensitive to changes than a longer one. Moreover, to reduce the computing time required to search for a change in the post-event detection window, its length should be as short as possible. The appropriate length of the post-event detection window was found to be 25~50% of the pre-change average window in order to get a relatively stable yet sensitive average for detection of on-off events. For the fan data sampled at one-minute intervals, the detection window was set to five data points, 50% of the pre-event averaging window. For the ASHRAE test building, the detection window was three data points, also 50% of the pre-event averaging window.

The Threshold for the Detection Statistic. Literature on the GLR method describes the detection threshold, a dimensionless quantity, as a trained parameter. The magnitude of an appropriate detection threshold scales with signal noise, the minimum signal change of interest (which comes from a knowledge of rated equipment-power levels), and the abruptness of potential changes in the system. For detection of on/off events from the total electrical power of eight fans in a campus building, the threshold was set to 200 on the basis of on/off tests. The same threshold was used at the ASHRAE test site. For general detection applications, this threshold might be set adaptively during tests designed to determine detection parameters appropriate for a given HVAC system. The threshold for the detection statistic can be initiated as an arbitrary small number, for example a magnitude of one, and then increased until all events of interest are identified with a minimum occurrence of false or missed alarms.

The Standard Deviation of the Power Data. The standard deviation is an important measure of data quality, which for HVAC-system electrical-power data may vary rapidly over time due to noise and changes in power of equipment of interest. In the GLR algorithm, the magnitude of the detection statistic, as shown in Equation 8 for the case where power changes of any magnitude are of interest, is inversely proportional to the standard deviation. Therefore, calculation of the standard deviation becomes a key issue for successful change detection. The simplest method to determine the standard deviation is to measure it on a one-time basis during a training period. However, GLR output based on a fixed standard deviation was rarely fully satisfactory, which prompted one of the improvements to be discussed in the next section.

The standard deviation is valuable information of itself, in addition to its impact on the GLR detection statistic. It will increase noticeably when a fan, pump, or chiller under closed-loop control is unstable due to poorly selected controller gains. Norford and Leeb (1996) showed that a poorly tuned chiller controller could be seen in the electrical signal measured at the HVAC service entry. Detection of the oscillations can be automated via calculation of the standard deviation in the data window.

IMPROVEMENTS TO THE DETECTION ALGORITHM

Resetting the Detector after an On/Off Event Has Been Detected. Consider a single on/off event, detected by the GLR algorithm as it works its way through a continuous stream of data, sliding the pre-event mean window over one data point at a time and then calculating the detection statistic for each time in the detection window. When the detection statistic exceeds the threshold, it is desirable that the alarm be immediately silenced. That is, a single event should produce a single needle spike in the detection statistic, because an alarm of appreciable width can mask subsequent events. However, close examination of the algorithm reveals that it can continue to produce alarms as the windows slide through the data. Only when the windows move entirely past the on/off event will the alarm necessarily cease. This phenomenon clearly depends on the length of the windows, the abruptness of the event, signal noise, and the value assigned to the threshold. While it is possible to adjust window width to account for this effect, a better strategy is to purge the pre-event mean window when an alarm first occurs and refill it with new data. At the time of purge, the detection statistic will immediately drop down below the threshold. Figure 2 shows a representative pre-event averaging window, detection window, and reset window, used to reload the pre-event averaging window. The length of the reset window is, of course, the same as the pre-event window.

Updated Standard Deviation. The standard deviation can be continuously calculated from data in the pre-event averaging window. While such an on-line calculation is more accurate than a single calculation during the GLR setup period, it was found necessary to limit the range of allowable values because such extremes unduly influence the magnitude of the detection statistic. For example, a very steady, noise-free electrical signal will have a low standard deviation and the detection statistic will be very large, even for events associated with very small power changes that are of no interest. A very noisy electrical-power signal will cause the detection statistic to be very small, potentially masking events of interest.

Limits on the standard deviation were assigned as a fraction of the total power. Tests showed that the standard deviation tends to increase as more equipment is in operation and the total power increases. Data from a training period were used to determine the ratio of the measured standard deviation and the measured total power. During subsequent on-line FDD tests, the standard deviation was estimated as a product of this

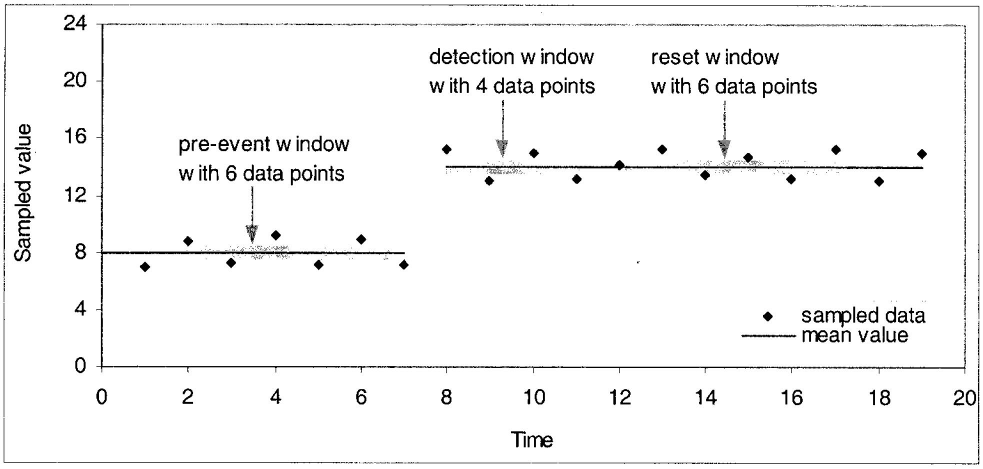

Figure 2 *Data windows used with the GLR detection algorithm.*

ratio and the measured total power. This method has given the most reliable estimate of the standard deviation because it eliminates the effect of the extreme values of the standard deviation while incorporating an updated estimate of the standard deviation in the calculation of the GLR detection statistic. At the ASHRAE test site, reasonable upper and lower limits for the standard deviation as fractions of the current power data were taken to be 10% and 1% respectively.

Non-Zero Minimum Expected Change. The GLR equation is easy to implement when the minimum expected change in power is zero. However, the zero-power minimum may cause more false alarms than a minimum value assigned on the basis of knowledge of equipment size. This can be readily seen from Equation 8: g_k increases with decreasing $\hat{V}_j$. In practice, it is often reasonable to find and set some minimum expected change based on the knowledge of the system and its components. For example, if the NILM were used to analyze fan performance after a variable-speed-drive retrofit, the minimum power level of interest would be set by the fans and would exclude smaller pieces of equipment, such as chilled-water pumps, as well as unknown disturbances. With the properly determined minimum expected change V_m, the detection statistic is insensitive to such disturbances and their accumulation within one window length, thus making the detection more reliable. A minimum expected change of 200 W was selected for the ASHRAE test building.

Median Filter. The performance of the GLR is adversely affected by signal noise. One way to reduce the impact of noise is to pre-process the data with a median filter (Karl et al. 1992), which simply picks the median value of the sequence to represent the current value. The length of the median filter's window should not be shorter than the duration of an electrical spike, observed in this study to be generally less than five seconds, and not longer than the interval between two consecutive events, selected as 20 seconds. The median filter is designed for the base sampling rate or interval, in the event multi-rate sampling is used. For example, if the base sampling interval is one second, then the number of data points in the filtering window should be between 5 and 20. The filtered data are then used for detection with different sampling rates, if necessary. At the ASHRAE test site, ten data points were included in the median filter window.

One possible problem with the median filter is that it masks rapid power oscillations. At the ASHRAE test building, this potential masking was avoided by using different data sets, one for change detection with the power data processed by a median filter and the other for oscillation detection without the median filter.

Combined Detection of On/Off Power Changes and Power Oscillations. Oscillation of power caused by unstable control in HVAC systems may degrade equipment and in some cases increase energy consumption, and is a fault that should be detected. One of the major characteristics of oscillations is the deviation of the data from their mean value in a window of appropriate length. This fault was identified by comparing the standard deviation of the data against a threshold that was dynamically adjusted as a fraction of the current power data.

The key points in designing a GLR detector with oscillation detection capability are proper thresholds for the detection statistic and for the standard deviation. Detecting step changes in power requires an upper threshold for the detection statistic, HGLR (high value for GLR), and a lower threshold for the standard deviation, LSTD (lower value for standard deviation). To detect an on-off event, the GLR detection statistic must exceed HGLR and the standard deviation must be no larger than LSTD. The oscillation detector relies on an upper threshold, HSTD, for the standard deviation and a lower threshold, LGLR, for power changes. The standard deviation must exceed HSTD and the GLR detection statistic must be no larger than LGLR. Thresholds established in this manner eliminate the fuzzy intermediate region where, for example, fluctuations in power may cause false alarms in the change detector, and permit a clearer distinction of step changes and

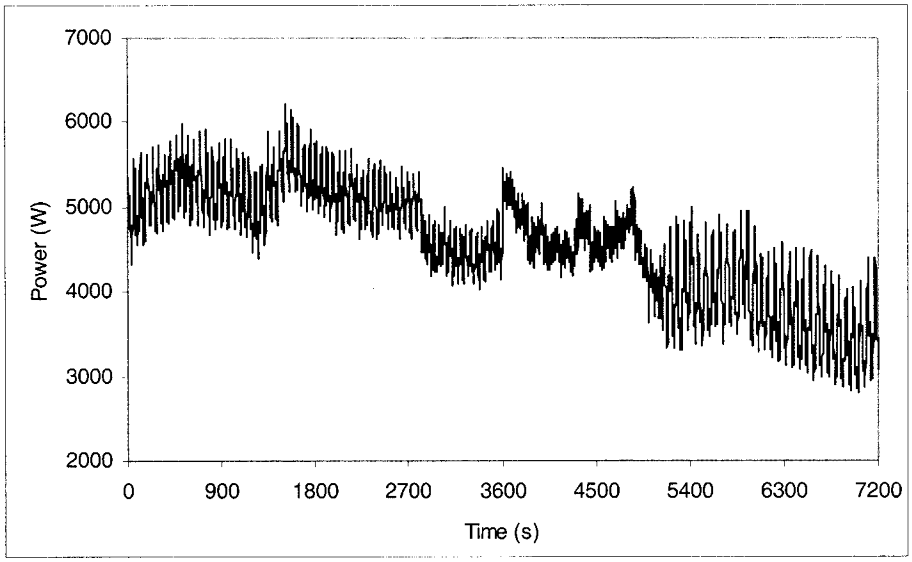

Figure 3 *Total electrical power of the HVAC system in a test building, showing oscillations indicative of an unstable supply-duct static-pressure controller as well as an on-off switching event.*

oscillations. The LGLR and the LSTD were set to be less than 10% of the HGLR and 20% of the HSTD, respectively, and were trained with measured data to minimize false alarms. Figure 3 shows data with both on-off events and an unstable controller, using power data recorded by a NILM installed at the motor-control center in the ASHRAE test building. During a period when the supply-fan static-pressure controller gain was set to produce power oscillations, a chilled-water pump was turned on at about 3600 seconds. With HGLR= 2500, LGLR= 250, HSTD= 0.02, and LSTD= 0.004, both events were alarmed successfully.

OPERATION OF THE GLR DETECTOR WITH MULTIPLE SAMPLING RATES

In spite of the improvements just described, applying the GLR detector to data sampled at a single rate inevitably involves a conflict between sensitivity to rapidly occurring events and susceptibility to false alarms generated by electrical noise. It is therefore worthwhile to explore the benefits of applying the GLR to a data set sampled at different rates. This will be done in three steps. First, the rules established above for determining the sampling rate will be applied and will be shown to have unavoidable limitations in detecting rapidly sequenced events. Second, the same rules will be shown to be unsatisfactory in detecting chiller on/off cycling in a noisy electrical signal. Third, the GLR detector will be applied at different sampling rates and the results combined in a way that overcomes some of these limitations.

Order of Magnitude Analysis of Sampling Rate Selection

Selecting a single sampling rate requires considerable care. An appropriate starting point is to bound the sampling interval: the upper bound should be less than the shortest time interval between two consecutive events of interest and the lower bound should be greater than the time required for the fastest turn-on event of interest. The need for the upper bound is clear, but it may lead to sampling rates that leave the GLR sensitive to noise, as will be discussed shortly. The lower bound is needed because shorter sampling intervals will make a step-change take on ramp-like properties and be harder for the GLR algorithm to detect.

The interval between consecutive events can be very short, as illustrated with a single day of whole-building data, taken at the test building used for ASHRAE 1020-RP and previously shown in Figure 1. The data were collected with a NILM and were sampled at 24 Hz in the first stage of the analysis. During a single 30-second period (12300-12330 seconds) three devices were turned on sequentially. (The identity of these devices is not important for this analysis and was not determined. The whole-building NILM was used primarily to detect the reciprocating chiller, and other events were not classified.) Reducing the sampling interval from 60 to 10, 1, and finally 0.125 seconds improved the resolution. Figure 4 shows this progression.

The three turn-on events were taken to be a single step change at a sampling interval of 60 seconds. With the sampling interval reduced to 10 seconds, the events were shown as single data points of different magnitude, discernible to the eye but not to the detection algorithm. When the sampling interval was reduced to 1 second (not shown in

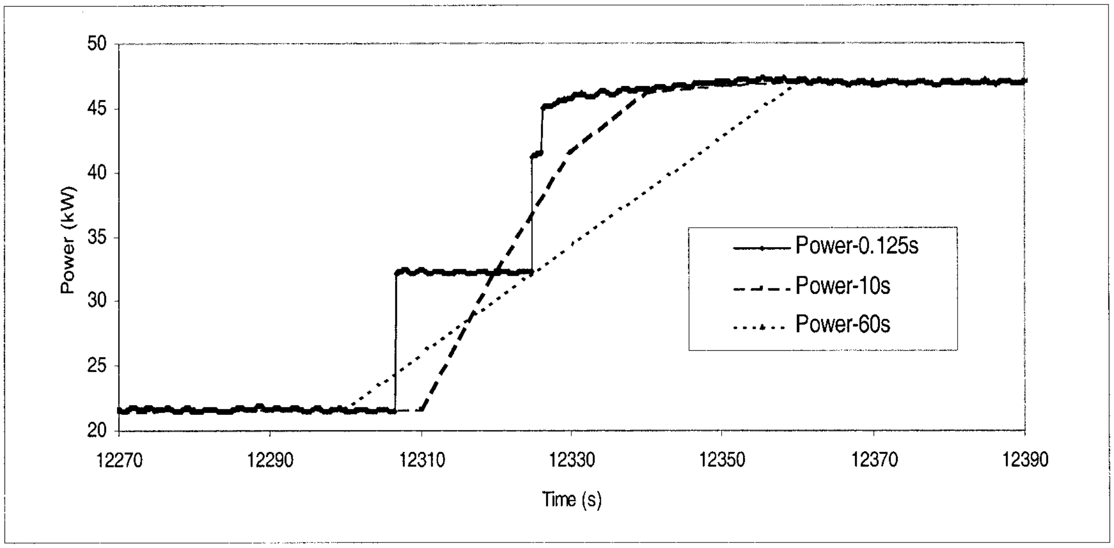

Figure 4 *Power data sampled at three different intervals during a period of 120 seconds when three events occurred, demonstrating the significantly different data patterns fed to the detector due to the varying sampling intervals.*

Figure 4 because it nearly overlapped data taken at 0.125 second intervals), the first and second events could be clearly recognized by eye and by the detection algorithm as well. The last event was discernible to the eye but not to the detection program, which needed an appropriate stream of data samples to build up the GLR detection statistic. The last event remained ambiguous to the program until the sampling rate was increased to 8 Hz, i.e., an interval of 0.125 seconds.

At the ASHRAE test building, the duration of the fastest turn-on event was 0.125 second and the shortest time interval between events was 1.25 seconds. As just shown, a sampling interval of 0.125 second was needed to produce enough data points to distinguish start-up transitions separated by 1.25 seconds. But subsequent tests to detect cycling of a reciprocating chiller showed that the false alarm rate increased drastically when the sampling interval dropped below one second. Because the three unknown start-up signals were not of interest and reliable detection of the chiller was central to the research, the minimum sampling interval was set to one second. Events spaced more closely than the sampling rate were unavoidably missed or misinterpreted. Such events seem to happen by coincidence and are not often seen in an HVAC system. In tests to date, there was generally an interval of not less than 10 seconds between on- and off-transitions for such coupled equipment as supply and return fans associated with the same air handler. The slower sampling rate not only reduced false alarms but also reduced the execution time of the GLR detector. This was important when the GLR detector was operated with multiple sampling rates as detailed later. The running time of the detector, as implemented on a personal computer with a clock speed of 233 MHz, was reduced from 90 minutes at an 8 Hz sampling rate to five minutes at a 1 Hz sampling rate for the evaluation of a single day's test, making the detector more desirable for on-line detection with comparable computers.

Detection of On/Off Changes with a Single Sampling Rate

Different sampling rates were used to examine power data for six days from the ASHRAE test building, with the goal of detecting reciprocating-chiller cycling. Figure 5 illustrates the first seven hours of the total power data on a single day. Figure 6 shows chiller power as measured with a submeter as well as GLR detection output for the chiller for a single sampling rate of 1 Hz. Note from Figure 5 that it is not easy to visually distinguish the chiller cycles from the whole-building power signal. The first chiller cycle, aligned with data from the submeter, is readily discerned. Most of the following events are obscured by other events of equal or greater magnitude.

The GLR detector identified all 13 start-up events from the 1 Hz samples but missed four of the shut-down events, which were masked by concurrent events. Other sampling intervals, including 2, 3, 4, 5, 6, 7, 8, 10, 12, 15, 20, 30, 40, 50, and 60 seconds, were also tested. However, all those sampling intervals produced results similar to those shown in Figure 6. With any one single sampling interval, the detector was not able to find all the on/off switches correctly. However, it was possible to visually identify all of the events by matching the on's and off's among outputs with different sampling rates. In this case, on/off matching among the outputs with the sampling rates of 1, 2, and 5 seconds identified all the on/off switches without any false alarms or missing events. Such other combinations as 1, 2, and 10 seconds were also successfully used for the matching. Similar detection patterns with single sampling rates were found with the remaining five

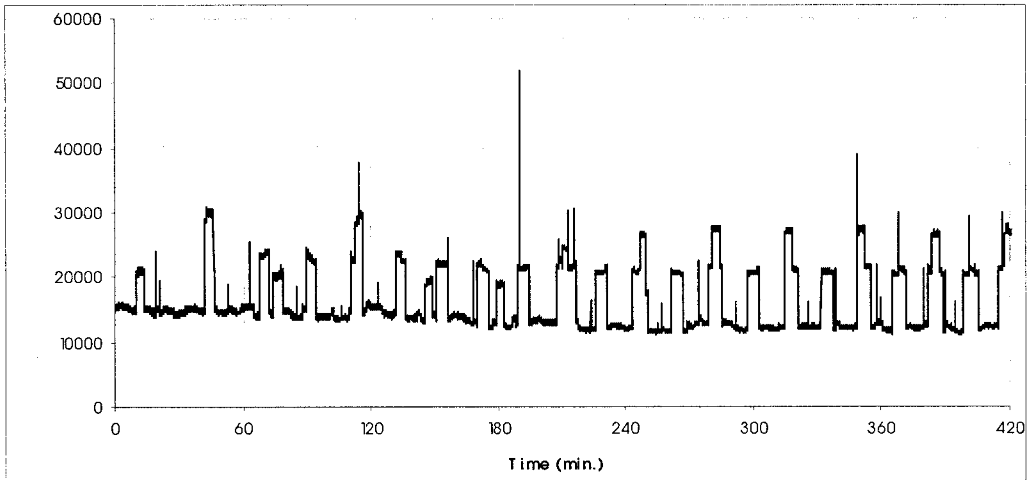

Figure 5 *Whole-building power data from the ASHRAE 1020-RP test site, sampled at 24 Hz and plotted at 10-second intervals. The time period is about seven hours.*

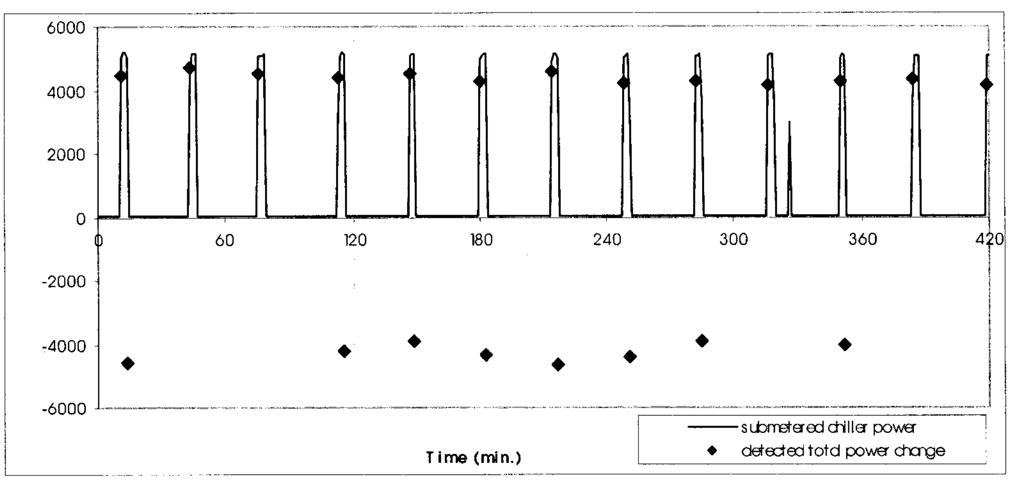

Figure 6 *Submetered chiller electrical power and the on/off cycles detected by the GLR algorithm as changes in the building's total power data sampled at 1 Hz over the same time period as in Figure 5.*

days' data. This prompted an effort to automate the process of matching events detected at different sampling rates.

Detection of On/Off Changes of Specific Equipment by Automatic Matching Among Multiple Sampling Intervals

The basic GLR detector was extended to analyze data at multiple sampling rates. Each sampled series with a specified integer-sampling interval between the lower and upper limits was supplied to the GLR detector. For the ASHRAE test building, an acceptable maximum value for the sampling rate was about an order of magnitude longer than the minimum value. Specifically, the minimum sampling interval was one second and the maximum was set to be 30 seconds.

Once an event was found by the detector, the power change was calculated as the difference between the mean in the post-event window and that in the pre-event window. The mean in the pre-event window is the average of data in the whole window, while that in the post-event window is calculated with the data from the point where the event is found to the end point of the window. The detected power changes for each sampling interval were sorted by the time of the changes, assigned to the equipment with that magnitude, and matched with the output from other sampling intervals. Matching of events involved the time of event and the sign (positive for on-transitions and negative for off-transitions) and the magnitude of the changes.

Upper and lower limits need to be trained for the magnitude of change or set adaptively (Hart 1992). Such limits can be trained by applying the detector to data collected during one day's operation. For the detection shown in Figure 6, for example, data during the first seven hours of a normal operation day will suffice for training. At some time point, if a positive change within the given limits is found, then the next negative change within the magnitude limits seen by any of the employed sampling rates will be assigned to the positive change and an on/off cycle is recorded.

Data filtering criteria may be applied to specific equipment to reduce the false alarm rate. This is helpful when the power of a component of interest is similar to that of other equipment, as is the case at the test building (Figure 5). For the chiller, for example, these criteria include the minimum off-time between cycles, which is typically set within the chiller controls to prevent unnecessary equipment cycling, and the minimum expected on-time. These limits were incorporated in the chiller-detection algorithm. Such filtering criteria are based on normal, not faulty, chiller operation. The above "minimum" criteria can be used for the detection and diagnosis of such faults as short cycles or prolonged operation time. This can be done with two different output files, one using the "minima" for filtering and the other using them for fault detection and diagnosis. If an excessive number of equipment cycles is detected in an unfiltered data stream, then there is either a fault or there are other components of comparable magnitude. If the power magnitudes of different equipment are distinct from each other, then these "minima" will not be needed for filtering and can be limited to detection of cycling faults.

The multiple-sampling-rate GLR approach has proved to be very useful in coping with electrical-power complexities not found in residential buildings. First, start-up and shut-down events vary in their apparent abruptness, as a function of the equipment (soft-start motors, for example) and as influenced by changes in other loads. In the ASHRAE test building, VAV air-handler fans have a very slow start-up signature because they are controlled by variable-speed drives. With the multi-rate sampler, these start-up transients can be detected, as was shown for the VAV fans in the test building. Detecting the fans required that the maximum sample interval be increased from 30 seconds to 10 minutes. At the other end of the sampling-rate spectrum, abrupt shut-down transients are difficult to detect when masked by gradual changes in power drawn by other equipment. In these cases, the fast samplers included in the multi-rate algorithm work best. Second, as shown in Figures 7 and 8, power oscillations characteristic of a poorly tuned controller were more reliably detected via analysis of data sets taken at multiple rates. This approach tends to mitigate the problem of picking a sampling rate appropriate for detection of oscillations at an unknown frequency.

SUMMARY OF TRAINING GUIDELINES

Guidelines for training the GLR detector, presented in detail above, are summarized as follows:

1. Record electrical power for the circuits monitored by the NILM for one day under typical operating conditions. The sampling rate should be between 1 and 10 Hz for common HVAC systems.
2. Locate the events from the abrupt changes in the total power data and estimate the fastest and the slowest events.
3. Determine the base sampling rate for detection. The base sampling rate will be used as the fastest sampling rate if multi-rate sampling is employed. Therefore, power data sampled at this rate should be used to discern by eye each event of interest.
4. Determine the window lengths. The length of the detection window should contain at least two data points. It should not be longer than the interval between two consecutive events and never shorter than a disturbance. The length of the pre-event window is 2~4 times longer than the detection window. The length of the post-event window is the same as that of the pre-event window.
5. Calculate and estimate the standard deviation of the power data as a fraction of the current total power data for the detection, f. The lower/upper limits of the standard deviation can be set at 1~2 magnitudes lower/higher than the estimated value f.
6. Estimate the threshold for the detection statistic. A reasonable base value for the threshold is $1/f^2$. The threshold can then be trained by adjusting this value until all events of interest can be seen by the detector, with a minimum number of false alarms.

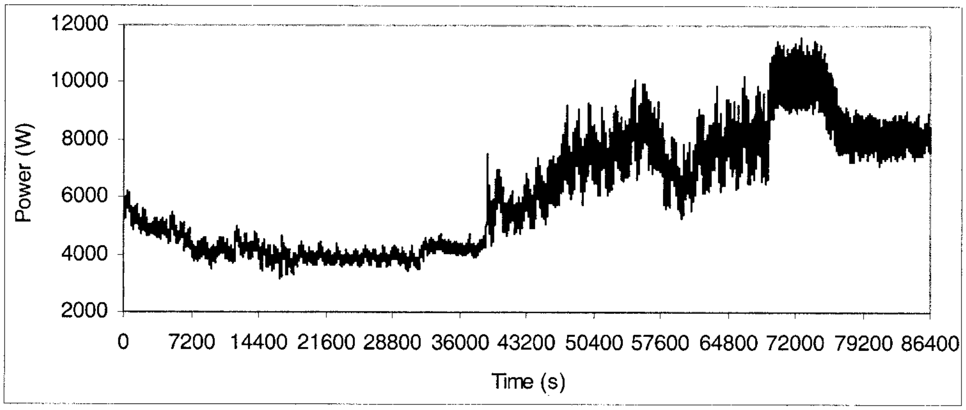

Figure 7 *Motor-control-center electrical-power data from the ASHRAE test site, taken over a 24-hour period beginning at 8 p.m. The controller gain for the supply fan in one of three VAV air handlers in the building (and monitored by the NILM on the motor-control center) was increased at about 7:25 a.m. to the point where the fan control was unstable.*

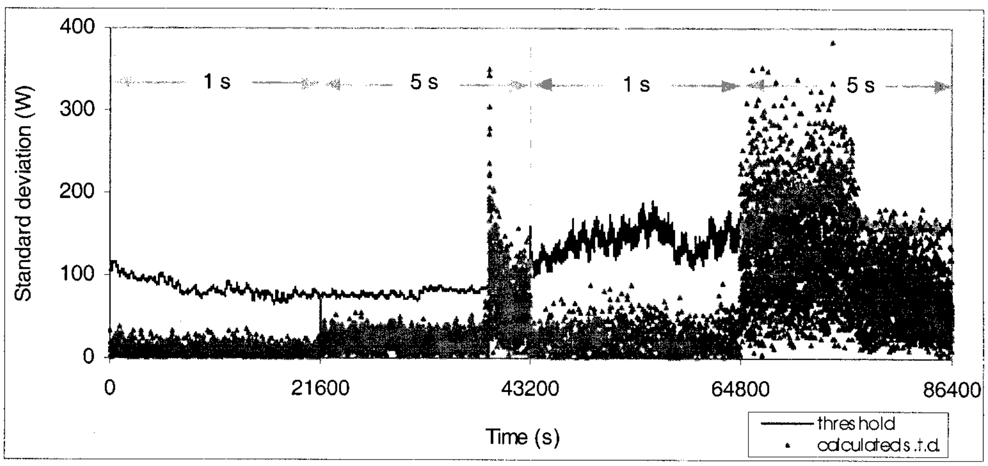

Figure 8 *Detection of the unstable supply-fan controller via calculation of the standard deviation with data sampled from the motor control center at two different intervals, 1 and 5 seconds, combined and shown in different time sections of 24-hour operation.*

It should be noted that exact values are not expected for the above parameters due to the statistical properties of the detection method. Slightly different combinations of these parameters may yield equally acceptable detection output.

The detector developed in this research has been successfully applied to another test site. With the above basic rules and guidelines, the training process for the parameters became much easier. Because the rules for the window lengths are for common HVAC systems, the same window lengths were used and the training was confined to determination of the thresholds for the detection statistic and the standard deviation. On/off events at this site were more abrupt than at the ASHRAE test building, in part because there were no variable-speed drives, and the detector produced satisfactory results while operating at a single sampling rate.

COMPARISON OF ON/OFF CHANGE DETECTION WITH CENTRALIZED AND SUBMETERED POWER DATA

A well-trained GLR detector will detect a high percentage of on/off events of interest, with a minimum number of false alarms. However, such performance does not necessarily mean that power changes can be quantified with sufficient precision to detect potentially faulty conditions. Table 1 and Figures 9-11 compare the GLR with single and multi-rate sampling against data from electrical submeters at the ASHRAE test building. During this five-hour test period, which started at 19:44 and included normal equipment operation and the end-of-day shut-down sequence, there were 17 on-off switching events for two hot-water pumps, two supply fans, and two return fans, among a total of six fans and 10 pumps served by the monitored circuits.

TABLE 1
ON-OFF POWER CHANGE DETECTION FOR A MOTOR CONTROL CENTER SERVING FANS AND PUMPS IN A TEST BUILDING, AS COMPARED WITH SUBMETERED DATA

Time	ON/OFF Equipment	Submeter (W)	NILM - GLR			
			Single Interval		Multiple Interval	
			Power (W)	Error (%)	Power (W)	Error (%)
20:10	Loop-B hot water pump—OFF	–405	–502	19.3	–400.6	1.1
20:19	Loop-B hot water pump—ON	264	249	6.0	284.1	7.1
20:21	Loop-A hot water pump—OFF	–275	Not found		–279.0	1.4
20:25	Loop-A hot water pump—ON	252	313	19.5	226.5	11.2
20:34	Loop-B hot water pump—OFF	–244	–98	149.0	–476.1	48.8
20:44	Loop-B hot water pump—ON	232	497	53.3	230.4	0.7
21:02	Loop-B hot water pump—OFF	–309	–346	10.7	-281.6	9.7
21:14	Loop-B hot water pump—ON	267	171	56.1	252.1	5.9
21:29	Loop-B hot water pump—OFF	–202	–11	1736.4	–171.5	17.8
21:47	Loop-B hot water pump—ON	237	289	18.0	257.2	7.9
21:55	AHU-A supply fan—OFF	–1400	–1396	0.3	–1410.9	0.8
21:56	AHU-A return fan—OFF	–82	Not found		Not found	
22:00	AHU-B supply fan—OFF AHU-B return fan—OFF	–430	–274	56.9	–485.1	11.4
22:06	Loop-B hot water pump—OFF	–300	–242	24.0	–243.2	23.4
22:26	Loop-A hot water pump—OFF	–264	–283	6.7	–244.4	8.0
23:40	Loop-A hot water pump—ON	342	373	8.3	220.1	55.3
00:33	Loop-B hot water pump—ON	295	494	40.3	450.5	34.5

With Single Sampling Interval

Of the 17 events, 15 were found with different levels of error and the remaining 2 were missed. Table 1 and Figure 10 show that the quality of the detection depended on the magnitude of the change relative to the total power as well as the current data trend. For example, when water pump B turned off at 20:34 and the change in power of –244 W was less than 5% of the total power, the error was about 150%, relative to the submetered change in power. However, when the supply fan for AHU-A turned off at 21:55 and the power change of -1400 W was about 25% of the total power, the error was less than 0.3%, again relative to the submetered power. When the change magnitude was very small, the detector was not able to find the event at all, including the missed event at 21:56 when the return fan for AHU-A was turned off with a power change of only –82 W.

The shorter the time interval between two changes, the more difficult it was to find the changes, especially small ones. This is simply because as a steady-state detector, the GLR needs some time for the effect of the former event to die out in order to find the subsequent change. The smaller the magnitude of the change, the longer the time interval needed to detect it. This is demonstrated by the output at 20:21 for the turnoff of pump A. With a magnitude of -275 W, the event was

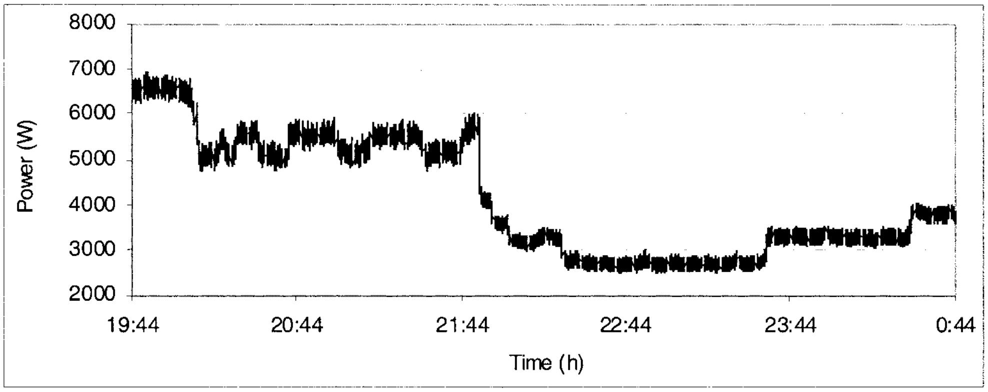

Figure 9 *Electrical power data collected at the motor-control center in the ASHRAE 1020-RP test building, taken over a five-hour interval beginning at 7:40 p.m. that includes the normal evening shutdown period. Seventeen on-off transitions are included in the data.*

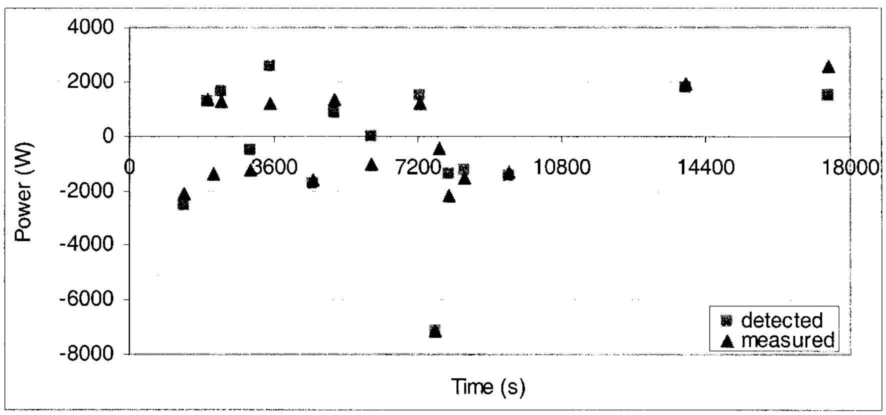

Figure 10 *Comparison of centralized and submetered power monitoring in the ASHRAE 1020-RP test building, with analysis of data taken at a single sampling rate of 0.1 Hz.*

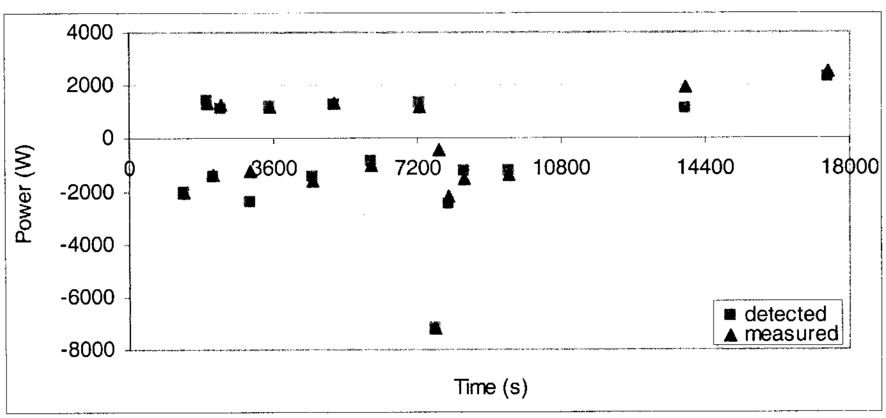

Figure 11 *Comparison of centralized and submetered power monitoring in the ASHRAE test building, with analysis of data taken at multiple sampling rates. Power monitoring is improved relative to analysis of data at a single sampling rate of 0.1 Hz.*

still missed because it was masked by gradual changes in electrical power due to the previous off-transition and to noise, which can be seen in Figure 9.

With Multiple Sampling Intervals

The multiple-sampling-rate approach has been shown to improve both the detectability of events and the resolution of the associated change in electrical power. This can be demonstrated by running the detector with multiple sampling intervals through the same five-hour data set used to evaluate the single-sampling-rate algorithm and plotted in Figure 9. As shown in Figure 11, the multiple-sampling-rate detector identified 16 of 17 on/off events, one of which was undetectable with the single-sampling-rate approach. Moreover, for almost all events the detector estimates a change in power that is closer to the submeter measurements than was the case for the single-sampling-rate algorithm, as listed in Table 1.

Under a few circumstances, the detection error for the multi-rate sampler was larger than that for the single sampling rate. With the lowered threshold of the detection statistic for multiple rates (the lowest value among the different sampling rates), the data pattern for reset of the detection window is different from any single sampling rate with a higher threshold, and this causes some variation of the change magnitude. But the occurrence of such problems is rare because the multi-rate detection is virtually a vote among different sampling intervals.

The major issue to be addressed in the near future for the multi-rate detection is to develop an appropriate algorithm to automatically calculate or select the "optimal" change value among different sampling intervals, i.e., the value that is closest to the real change. This might be realized by averaging, voting, or sorting among given sampling intervals on the basis of observations with more testing data.

DESCRIPTION OF NILM HARDWARE

A block diagram of the non-intrusive load monitor installation at the test facility used in ASHRAE RP-1020 appears in Figure 12. In addition to the conventional NILM-style connections at the main panel (for remote1) and the motor control center (for remote2), selected loads associated with air-handling unit B and the chiller were individually measured. With the connections shown in Figure 12, remote1 measured spectral envelopes (i.e., power envelopes for the fundamental and higher harmonic signals) for the entire building and remote2 measured spectral envelopes for the motor control center. Remote1 and remote2 shared the task of collecting data from individual loads.

The hardware platform for the prototype event detector consisted of a personal computer working in tandem with a digital signal processor (DSP). Custom software was written for the circuit board that included the DSP to compute estimates of the spectral or Fourier components of observed current waveforms in real time. The board provided the DSP chip, a 16-bit, 2-channel analog-to-digital converter (ADC), a 16-bit, 2-channel digital-to-analog converter (DAC), off-chip memory for temporary storage, an off-chip flash PROM for permanent storage, and a convenient program-debugging interface. The DSP communicated data to the host PC over a parallel port. The host PC processed the spectral information to provide load recognition and other value-added services such as critical load diagnostics. Because this prototype was PC-based, it was possible to provide these services and deliver information remotely.

The computers labeled remote 1 and remote 2 in Figure 12 operated at 200 MHz and were equipped with 32 MB memory and approximately 1 GB hard drives. Both machines were equipped with a network interface card, a data acquisition card, and the DSP card. The computers ran a highly stable, open-source, free operating system and were remotely maintained by secure shell over the Internet. Stability is a key consideration for any system that must operate remotely in the field. At the test site, both systems operated without any direct intervention for over two years of constant use. During active testing, data collected over 24-hour periods were accumulated and compressed at the on-site computers and were automatically transmitted at night to laboratory computers, where NILM data were maintained on an FTP site for retrieval and analysis.

The NILM was conceived as a means of reducing the cost of obtaining electrical-power data, and it is worth making a rough assessment of costs and benefits relative to currently available metering technologies. The costs can be compared with conventional AC watt transducers, which require a voltage tap and current transducers on the input side and produce as output a low-voltage or current signal proportional to power. Both the NILM and a watt transducer require a personal computer to collect data. This computer can stand alone as an independent data logger or can be the same personal computer used for an energy management system. Both the NILM and a watt transducer require an A/D board to digitize analog information for the PC. This information is voltage and current data for the NILM and the power data for the watt transducer. Both the NILM and the watt transducer require current transducers and a voltage tap. The NILM can use the same PC, A/D board, current transducers, and voltage tap as the watt transducer. The NILM eliminates the watt transducer itself, performing the convolution of current and voltage in software, but requires an inexpensive voltage transducer ($25 for a research-grade transducer and a fraction of that for a transducer of adequate bandwidth) to lower the line voltage to a level appropriate for the A/D converter. The cost of the NILM is therefore comparable to a submeter, even if the goal is to monitor a single load. Multiple loads require multiple watt transducers, each with current transducers. A single-phase watt transducer with current transducer costs about $250-300. There were six loads of prime interest at the motor-control center in the ASHRAE test building: three supply fans and three chilled-water pumps. The costs are

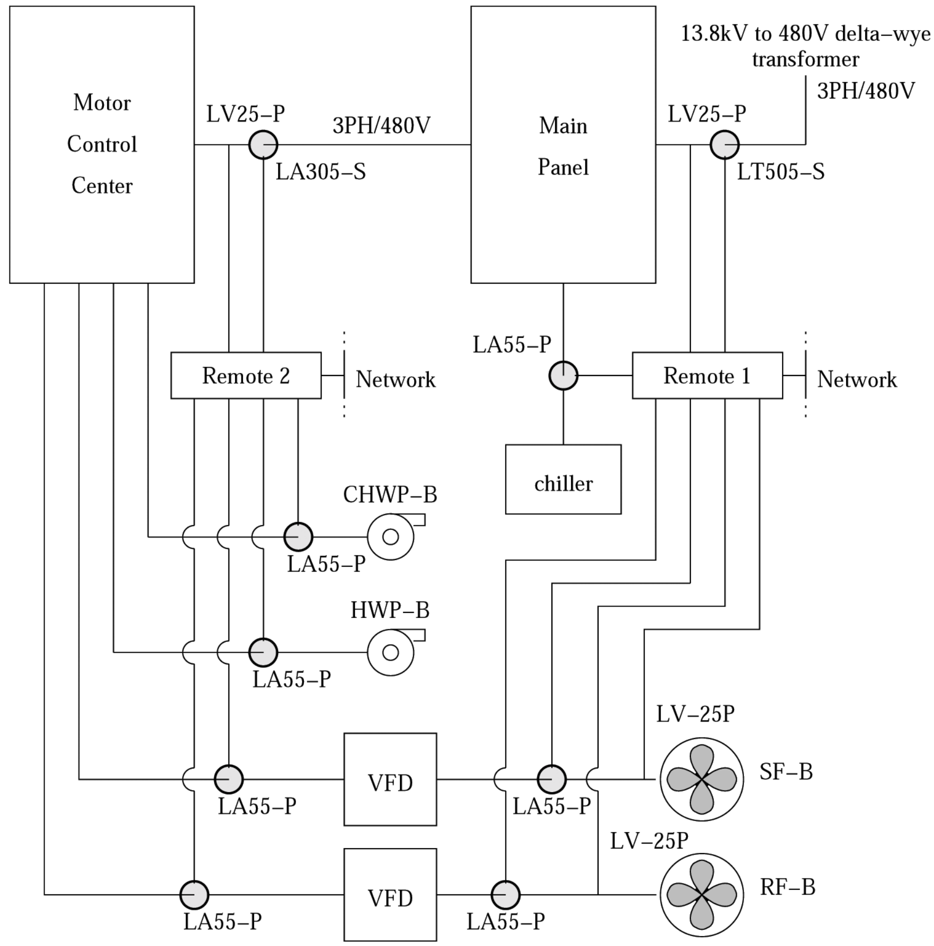

Figure 12 *Electrical schematic showing location of the two NILM meters at the test building. Sensors labeled LA55-P, LA305-P, and LT505-S are 50A, 300A, and 500A Hall-effect current sensors. LV25-P is an isolated voltage sensor.*

comparable for the first fan and the NILM saves the $1250-1500 required for the remaining two fans and three pumps.

More testing is required to determine whether the NILM can consistently produce power data of accuracy suitable for its intended use. Until these tests are done, submetered data remain the standard. But the NILM should not be evaluated as a trade-off between cost and measurement accuracy. It is a flexible platform that generates a data stream far richer than a submeter produces and that can analyze that data in ways limited only by the imagination of an engineer and skill of a programmer. For example, the power harmonics it calculates can be used for power-quality assessments and for identifying loads that generate certain harmonics. It can analyze very rapid start-up transients (Shaw et al. 2002) as a means of detecting faults in equipment performance and can detect power oscillations that are associated with actuator wear.

CONCLUSIONS

In principle, low-cost information about individual electrical components in a building can be obtained via careful analysis of power measurements at central locations within the building, notably the electrical service entrance and motor-control centers that supply power to HVAC components. Visual analysis of electrical power sampled at appropriate speeds shows step changes that can be associated with equipment of interest. The development of a reliable automatic detector suitable for use in noisy and complex electrical environments has involved selection of the basic statistical approach, development of guidelines for tuning the detector, and innovations that improve performance (enhanced sensitivity to signals of interest and rejection of electrical noise). Further work has extended the detection approach to identify sustained power oscillations, indicative of poorly tuned controllers, as well as step changes, and to operate on data sampled at multiple rates, in order to fill in missing events and reject false alarms that may be generated from a single data set. Test results show that the centralized or non-intrusive load monitor can detect all on-off events of interest in some data sets from real buildings and that further work is required to automate the process of tuning detector parameters. With other data sets, the detector detected most, but not all, events. Further testing in other buildings is required, particularly to evaluate and automate the steps needed to tune the detector. In recent years of development, the detector's capabilities have been enhanced while the hardware cost has decreased, encouraging future development.

ACKNOWLEDGMENT

The authors warmly acknowledge the financial and technical support of Honeywell, Inc., the California Energy Commission (via subcontracts from Lawrence Berkeley National Laboratory and Architectural Energy Corporation), and ASHRAE.

REFERENCES

Abler, C., R. Lepard, S. Shaw, D. Luo, S. Leeb, and L. Norford. 1998. Instrumentation for high-performance nonintrusive electrical load monitoring. *ASME J. Solar Energy Engineering*.

Basseville, M., and I. Nikiforov. 1993. *Detection of abrupt changes theory and application*. Prentice Hall Information and System Sciences Series. Englewood Cliffs, N.J.: Prentice Hall.

Deschizeau, M., P. Bertrand, A. Anglade, and M. Grimaldi. 2000. A new method for detailed electric consumption of domestic appliances. *Proceedings of 2d international conference on energy efficiency in households appliances and lighting, Naples, Italy.*

Hart, G.W. 1992. Nonintrusive appliance load monitoring. *Proceedings of the IEEE* 80(12): 1870-1891.

Hill, R.O. 1995. Applied change of mean detection techniques for HVAC fault detection and diagnosis of power monitoring. M.S. thesis, MIT Department of Architecture.

Karl, W.C., S.B. Leeb, L.A. Jones, J.L. Kirtley, Jr., and G.C. Verghese. 1992. Applications of rank-based filters in power electronics. *IEEE Transactions on Power Electronics* 7(3): 437-444.

Leeb, S.B. 1993. A conjoint pattern recognition approach to nonintrusive load monitoring. Ph.D. thesis, MIT Department of Electrical Engineering and Computer Science.

Leeb, S.B., S.R. Shaw, and J.L. Kirtley. 1995. Transient event detection in spectral envelope estimates for nonintrusive load monitoring. *IEEE Transactions on Power Delivery.*

Leeb, S.B., and J.L. Kirtley. 1996. A transient event detector for nonintrusive load monitoring. Technical report. U.S. Patent 5,483,153.

Leeb, S.B., U.A. Kahn, and S.R. Shaw. 1998. Multiprocessing transient event detector for use in a nonintrusive electrical load monitoring system. Technical report. U.S. Patent 5,717,325.

Norford, L.K., and N. Mabey. 1992. Nonintrusive electric load monitoring in commercial buildings. *Proceedings of the eighth symposium on improving building systems in hot and humid climates*. Dallas, Texas.

Norford, L.K., and S.B. Leeb. 1996. Nonintrusive electrical load monitoring. *Energy and Buildings* 24: 51-64.

Norford, L.K., J.A. Wright, R. Buswell, D. Luo, S.R. Shaw, and S.B. Leeb. 2000. Demonstration of fault detection and diagnosis methods in a real building, ASHRAE RP-1020.

Norford, L.K., J.A. Wright, R. Buswell, D. Luo, S.R. Shaw, and S.B. Leeb. 2002. Demonstration of fault detection and diagnosis methods in air-handling units, ASHRAE RP-102. To be published in *International Journal of HVAC&R Research*.

Rice, J.A. 1988. *Mathematical statistics and data analysis*. Pacific Grove, Calif.: Wadsworth & Brooks/Cold Advanced Books & Software.

Shaw, S.R., D. Luo, L.K. Norford and S.B. Leeb. 2002. Fault detection via electrical load monitoring. To be published in *International Journal of HVAC&R Research*.

DISCUSSION

Rick Danks, Masa Glenn Research Center, Cleveland, Ohio: 1) Was the power distribution checked for integrity to avoid false readings due to loose connections, etc.? 2) Can this system distinguish between system problems (some listed above) and equipment problems?

Leslie K. Norford: The power-distribution system in the test building was very well maintained and of recent vintage. It was inspected visually during the installation of the two centrally located power monitors but was not subjected to other tests. The monitors are capable of detecting a side variety of problems but must be programmed accordingly. A reasonable analogy is that of a spot light that must be pointed in the right direction to reveal an object of interest, in contrast with a flood light that requires less tuning but reveals less detail. The monitors can be used to detect and diagnose such system problems as voltage sags due to start-up of large components, which may cause other equipment to abruptly shut down. With one-time measurements of impedances in the building electrical system, the monitor can estimate the voltage distortion at any point in the building. At the equipment level, the monitor can detect faulty electrical contactors.

Grant Wichenko, Appin Associates, Winnipeg, Canada: 1) Have you used power quality monitoring? 2) What can be determined with you system that you cannot get with a normal DDC system? 3) What resolution do you need to make sense of the data?

Norford: The electrical monitoring system we describe takes current and voltage measurements at high speed, calculates harmonics of real and reactive power, and is capable of power-quality monitoring. A normal Building Energy Management System (BEMS) is not used for detection and diagnosis of HVAC faults. Extensive recent research, inside and outside ASHRAE, is producing FDD techniques that could be included in a BEMS. Our approach makes use of electrical-power data and could be used in lieu of or in addition to approaches that make use of thermofluid data (temperatures and flows). Further, our system is capable of finding some faults, such as misalignment in a motor-belt fan system, that thermal modeling would probably not detect.

Equipment-specific energy consumption is also important to building operators. In a conventional system, it is necessary to install expensive watt meters on each component of interest, an approach that is rarely taken due to cost. We are conducting field tests to determine whether we can provide an acceptable estimate of equipment-specific power via high-speed monitoring from a single point, a less expensive alternative.

The resolution required to make sense of the data depends on the application. We sample the current 128 times per line cycle, or 7,680 Hz. From these current data, we calculate real and reactive power and the fundamental and higher harmonics and produce a data stream at 120 Hz. We have down-sampled these data to obtain data at 1-10 Hz for FDD, based on step-changes in electrical power that indicate whether or not a pump, fan, or chiller has turned on or off. More fruitfully, we use the higher-speed data stream to resolve the start-up transients of motor-driven loads. These transients are very short, on the order of a second.