HVAC&R Research

ISSN: 1078-9669 (Print) 1938-5587 (Online) Journal homepage: http://www.tandfonline.com/loi/uhvc20

Detection and Diagnosis of HVAC Faults via Electrical Load Monitoring

S. R. Shaw , L. K. Norford , D. Luo & S. B. Leeb

To cite this article: S. R. Shaw , L. K. Norford , D. Luo & S. B. Leeb (2002) Detection and Diagnosis of HVAC Faults via Electrical Load Monitoring, HVAC&R Research, 8:1, 13-40

To link to this article: http://dx.doi.org/10.1080/10789669.2002.10391288

Published online: 03 Mar 2011.

Submit your article to this journal

Article views: 116

View related articles

Citing articles: 16 View citing articles

Full Terms & Conditions of access and use can be found at
http://www.tandfonline.com/action/journalInformation?journalCode=uhvc20

Download by: [18.62.19.151] **Date:** 19 July 2016, At: 07:55

Detection and Diagnosis of HVAC Faults via Electrical Load Monitoring

S.R. Shaw, Ph.D. **L.K. Norford, Ph.D.** **D. Luo, Ph.D.** **S.B. Leeb, Ph.D.**
Member ASHRAE Student Member ASHRAE

Detection and diagnosis of faults (FDD) in HVAC equipment have typically relied on measurements of variables available to a control system, including temperatures, flows, pressures, and actuator control signals. Electrical power at the level of a fan, pump, or chiller has been generally ignored because power meters are rarely installed at individual loads. This paper presents two techniques for using electrical power data for detecting and diagnosing a number of faults in air-handling units. The results from the two techniques are compared and the situation for which each is applicable is assessed.

One technique relies on gray-box correlations of electrical power with such exogenous variables as airflow or motor speed. This technique has been implemented with short-term average electrical power measured by dedicated submeters. With somewhat reduced resolution, it has also been implemented with a high-speed, centralized power meter that provides component-specific power information via analysis of the step changes in power that occur when a given device turns on or off. This technique was developed to detect and diagnose a limited number of air handler faults and is shown to work well with data taken from a test building. A detailed evaluation of the method is presented in the companion paper, which documents the results of a series of semiblind tests.

The second technique relies on physical models of the electromechanical dynamics that occur immediately after a motor is turned on. This technique has been demonstrated with submetered data for a pump and for a fan. Tests showed that several faults could be successfully detected from motor startup data alone. While the method relies solely on generally stable and accurate voltage and current sensors, thereby avoiding problems with flow and temperature sensors used in other fault detection methods, it requires electrical data taken directly at the motor, downstream of variable-speed drives, where current sensors would not be installed for control or load-monitoring purposes.

Downloaded by [] at 07:55 19 July 2016

INTRODUCTION

The performance of many HVAC systems is limited more by poor installation, commissioning and maintenance than by poor design (Tong 1989). Computer-based control systems have the capability to collect and store sensor and control signals that could be analyzed to detect and diagnose faults. A considerable amount of research work has been carried out to develop FDD techniques for HVAC systems and, recently, to test these techniques in realistic laboratory settings and in real buildings (Ahn et al. 2001; Chen and Braun 2001; Dexter and Benouarets 1996; Dexter and Ngo 2001; Haves et al. 1996a; House et al. 2001; Hyvarinen 1996; Lee et al. 1996a, 1996b; Li et al. 1996; Peitsman and Bakker 1996; Salsbury 1996; Stylianou and Nikanpour 1996; Tsutsui and Kamimura 1996; Yoshida et al. 1996).

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

A fault is detected when the observed behavior of a system differs from the expected behavior by some threshold. The expected behavior of the system is often expressed in a model, whether physical, statistical or fuzzy. The number of measurements required to use a given model for fault detection is of central importance, as there is a cost associated with installation and maintenance of sensors that are not required by the control system.

A fault is diagnosed when it is detected and a cause is determined. Diagnosis is significantly more difficult than detection, because measured effects must be attributed to a particular cause to the exclusion of all other possibilities. The difficulty of diagnosis increases with the number of possible causes, and careful validation in the context of the target application is required for any proposed diagnostic technique.

Current approaches to HVAC FDD have been driven by the confidence researchers have in their own modeling approach—its ease of implementation, robustness, and ability to generalize to many different HVAC faults—and in the availability and accuracy of measured data. The approach in this paper begins with the premise that electrical power measurements are useful in FDD, based on the authors' experience extracting device-level information from centralized power measurements. The paper presents FDD techniques using submetered power data that, in some cases, may be readily adapted for use with centralized measurements. Whether these techniques are economically and technically useful needs to be explored in the field over the coming years.

Downloaded by [] at 07:55 19 July 2016

There are four generic approaches to consider for use of electrical power data in HVAC FDD, as shown in Table 1. Power signals can be obtained from meters attached to fans, pumps, chillers, or other individual pieces of electrically powered HVAC equipment. Alternatively, power data can be collected at a single point. Further, detection and diagnosis of faults can be based on analysis of changes in steady-state electrical power, or analysis of the dynamic variation of power during the typically very brief startup transient. For submetered power, steady-state is defined as one-minute averages at times other than those encompassing a startup or shutdown event or substantial power oscillation. In contrast, for centralized power measurements, steady-state power is the difference in power before and after a turn-on or turn-off event. Techniques presented in this paper cover three of the four approaches, as noted in Table 1. The detection and diagnosis of faults from startup transients recorded centrally, the most powerful but difficult method, has not been attempted to date.

Centralized load monitoring takes its name of Non-Intrusive Load Monitoring (NILM) from its origins in load analysis in houses, where the normal revenue meter was replaced by a meter capable of clustering step changes in real and reactive power and associating these clusters with major household appliances (Hart 1992). Over the last decade, the early techniques have been extended in two directions. First, enhanced detection of step changes in power have been shown to screen out power spikes caused by switching electronics and to find power oscillations caused by poorly tuned controllers (Norford and Mabey 1992, Hill 1995, Luo et al. 2002). Second, a powerful detection approach has been developed to analyze the rapid changes in power that occur when a motor, lamp, or computer is first turned on (Leeb 1993, Leeb et al. 1995, Norford and Leeb 1996, Abler et al. 1998). The startup patterns are governed by the physics of the device and are generally not masked by power electronics that reduce reactive power. In addition, anal-

Table 1. Options for Detecting and Diagnosing HVAC Faults by Analyzing Electrical Power Signals Explored in this Paper

Location of Meters	Analysis of Changes in Steady-State Power	Analysis of Power Dynamics During Equipment Startup
Individual loads	X	X
Centralized	X	

ysis of startup transients is useful in a busy electrical environment, due to the very brief amount of time (typically fractions of a second) required to log a characteristic signal. Both approaches to centralized load monitoring rely on a combined hardware/software platform that is considerably more capable than the typical Watt meter and of comparable cost. This platform has been refined to the extent that it consists solely of the necessary transducers fed into a personal computer for analysis and communication via the internet.

This paper focuses on the use of electrical power signals for HVAC FDD and not on centralized load monitoring. One major motivation for this work has been an ASHRAE-sponsored research project, 1020-RP, "Demonstration of Fault Detection and Diagnosis Methods in a Real Building." In this project, one electrical power FDD approach, using gray-box correlations of steady-state changes in electrical power with flow or other variables, was tested on three variable air volume (VAV) air-handling units (AHUs) in a heavily instrumented test building. A second FDD approach (not required for 1020-RP) was also tested with data from the test building. This approach uses a physical model of the dynamics of a motor and driven load during a startup transient. Parameters estimated solely from electrical current measurements are compared with parameters estimated during periods of normal behavior to detect faults.

The HVAC system, diagrammed and described more fully in the companion paper, Norford et al. (2002), includes variable-speed supply and return fans, constant supply-air and chilled water temperature control, and a primary-secondary pumping system with constant-speed secondary pumps. The seven AHU faults included in these tests, listed in Table 2, are featured in this paper. Table 2 notes whether the selected faults were abrupt or occurred slowly over time; as implemented in the test building, abrupt faults were introduced as such and degradation faults were introduced over one- to three-day periods. Electrical power FDD methods are no different from others in their ability to find abrupt faults more easily than degradations.

Test results presented in this paper are limited to the faults introduced in the test building and as such are demonstrations of the methods rather than comprehensive assessments of their efficacy. A final report and the companion paper (Norford et al. 2000, 2002) summarize the results of the blind tests conducted as part of 1020-RP. This paper, in effect, lays a foundation for the summary paper.

While the presentation focuses on a small number of artificially introduced faults, the presented FDD methods can in principle be extended to cover additional AHU faults and faults in other systems. The obvious prerequisite is that any fault to be detected by these methods must

Downloaded by [] at 07:55 19 July 2016

Table 2. List of Air-Handling Unit Faults Detected and Diagnosed with Electrical Power Data

Fault	Type
Air Mixing Section	
Stuck-closed recirculation damper	Abrupt
Leaking recirculation damper	Degradation
Filter-Coil Section	
Leaking cooling coil valve	Degradation
Reduced coil capacity (water-side)	Degradation
Fan	
Drifting pressure sensor	Degradation
Unstable supply fan controller	Abrupt
Slipping supply fan belt	Degradation

Note: Fault implementation is described in Norford et al. (2002).

cause a change in electrical power. Further, the power change must be sufficiently large to cause a noticeable deviation from expected behavior, which in turn must be defined on the basis of gray-box power correlations using steady-state data or a physical model of a motor's startup transient. The gray-box power correlations require an understanding of the physics of the system under investigation, to determine power correlations suitable for detecting a given fault. The method therefore requires a considerable amount of expert knowledge to be made more generally applicable.

APPROACH

A. Detection and Diagnosis of HVAC Faults with Gray-Box Models and Submetered Electrical Power Measurements Under Steady-State Conditions

The gray-box electrical power method for fault detection and diagnosis rests on the following four steps:

1. A training phase, consisting of correlation of electrical power as measured under steady-state conditions with load measurements (gray-box models)
2. Measurement of electrical power
3. Detection of faults by comparison of measured and predicted electrical power and by screening for rapid power oscillations
4. Diagnosis of faults by rule-based analysis of power deviations

The gray-box method requires that fan, pump, and chiller power be correlated with an indicator of load for each electrical component. These correlations establish models for detecting equipment faults. For equipment such as that installed in the test building, Tables 3 and 4 summarize how these correlations depend on building thermal loads, as expressed by outside temperature in each of four regions:

1. $T_{out} < T_{balance\ point}$
2. $T_{balance\ point} < T_{out} < T_{supply\ air} - \Delta T_{supply\ fan}$
3. $T_{supply\ air} - \Delta T_{supply\ fan} < T_{out} < T_{return\ air}$
4. $T_{return\ air} < T_{out}$

The balance point temperature is the outside dry-bulb temperature at which the building requires neither heating nor cooling. In commercial buildings, $T_{balance\ point}$ accounts for a minimum intake of outdoor air required to satisfy ASHRAE *Standard* 62. Above $T_{balance\ point}$ there is a net cooling load and the supply airflow will increase. In some cases, $T_{balance\ point}$ may exceed $T_{supply\ air}$, in which case the two are simply exchanged in the temperature segments. The supply air temperature is assumed to be constant, as was the case in the test building. In Table 4, the coil capacity and slipping fan belt faults are best detected at high cooling loads but can be detected at lower loads if sufficiently severe.

Correlations of Fan Power with Airflow

If electrical power measurements for HVAC equipment are used as the basis for an FDD method, there must be some model of power under known load conditions. Studies have established that whole-building fuel use correlates with outdoor temperature (Fels 1986, Ruch and Claridge 1992). At the device level, fan power is of key importance in detecting many AHU faults. Previous work (Lorenzetti and Norford 1992) showed that hourly average power measurements for VAV supply and return fans correlate well with outdoor dry-bulb temperature. A simple three-variable change-point model was found to work as well as or better than other

Downloaded by [] at 07:55 19 July 2016

Table 3. Variation of Outside and Supply Airflows and Pump, Chiller, and Fan Powers with T_{out}[a]

Temperature Region and Control Mode	Outside Airflow	Supply Airflow	Pump Power	Chiller Power	Fan Power
1 Heating	Constant at minimum value	Constant at minimum value	Off	Off	Constant at minimum value
2 Free cooling via air-side economizer	Increasing with T_{out}	Increasing with T_{out}	On; constant power	Energized; minimal power to meet thermal loads in piping	Increasing as a polynomial function of flow or speed
3 Mechanical cooling with 100% outside airflow	Increasing with T_{out}	Increasing with T_{out}	Nearly constant power if piping is balanced	Increasing with T_{out}	Increasing as a polynomial function of flow or speed
4 Mechanical cooling with minimum outside airflow	Constant at minimum value	Increasing with T_{out}	Nearly constant power if piping is balanced	Increasing with T_{out}	Increasing as a polynomial function of flow or speed

[a]For a VAV system with variable-speed supply and return fans, constant supply air temperature, and a constant-speed secondary chilled water pump.

Table 4. Detectability of Faults with Electrical Power Data

Temperature Region	Stuck-Closed Recirculation Damper	Leaky Recirculation Damper	Leaky Cooling Coil Valve	Reduced Cooling Coil Capacity	Pressure Sensor Drift	Unstable Pressure Controller	Slipping Supply Fan Belt
1	Yes	No	No	No	Yes	Yes	Maybe
2	No	No	Yes	No	Yes	Yes	Maybe
3	No	Yes	No	Maybe	Yes	Yes	Maybe
4	Yes	No	No	Yes	Yes	Yes	Yes

Downloaded by [] at 07:55 19 July 2016

models fitted to these data. Outdoor dry-bulb temperature is a reasonable predictor because the sensible fraction of envelope loads (latent heat transfer excluded) influences the total amount of air delivered to occupied spaces. However, this predictor does not account for variations in airflow, and hence fan power, due to changes in internal or solar loads. Correlations with measured airflow provide a more precise estimate of fan power, as established in an earlier study of applications of electrical load monitoring to fault detection in ventilation systems (Norford and Little 1993). Such correlations have also been used to estimate fan energy consumption before and after variable-speed drive retrofits (Englander and Norford 1992).

Figure 1 shows a third-order polynomial correlation between fan power and airflow for a VAV supply fan with a variable-speed drive. The use of a third-order polynomial correlation is based on the fan laws, which show that power varies as a cubic function of speed for a variable-speed centrifugal fan. A similar correlation has also provided a good fit in practice for data collected from VAV fans equipped with inlet vanes. Ninety percent confidence intervals were established from uncertainties in the polynomial coefficients and a *t*-statistic. The confidence intervals express the confidence that a single new measurement point will lie between the upper and lower intervals, if the measurement is subject to the same conditions as occurred during the training phase. Increasing the confidence interval would make the method less sensitive to faults and less likely to generate false alarms. In the test building, the use of 90% confidence intervals did not generate false alarms, but in practice the confidence interval could be enlarged, reducing the number of both detected faults and false alarms. To tighten the correlation and improve the sensitivity of the method, only training data with duct static pressure within 25 Pa (0.1 in. of water) of the 300 Pa (1.2 in. of water) set point were accepted. This tolerance on duct static pressure was arbitrarily selected after the data were examined by eye and was intended to elimi-

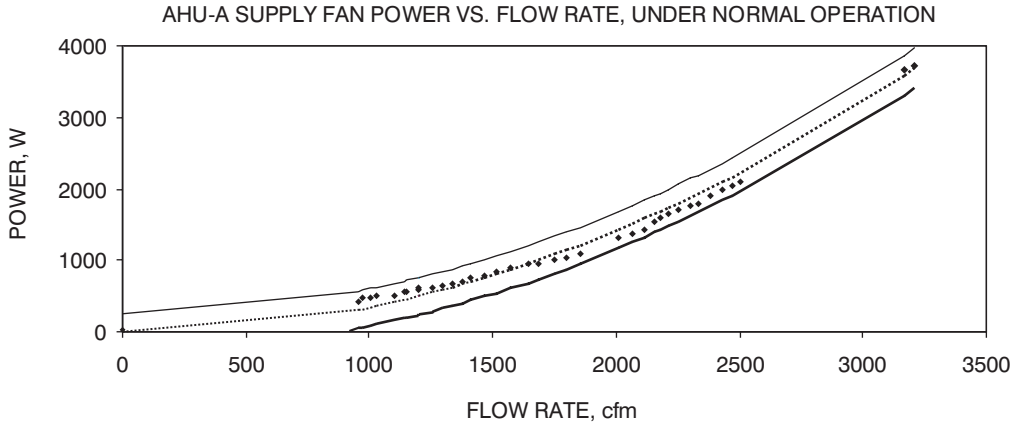

Figure 1. Correlation of Fan Power with Airflow During a Training Period

Downloaded by [] at 07:55 19 July 2016

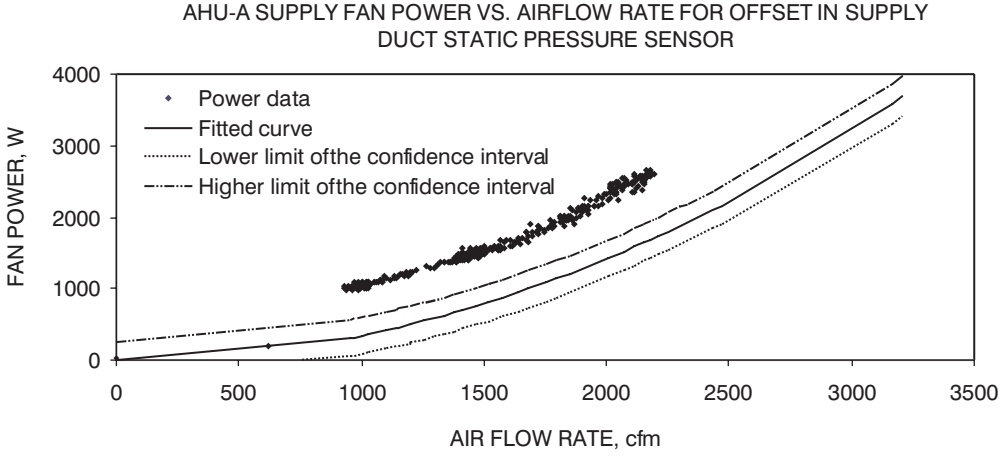

Figure 2. Fan Power Data Outside 90% Confidence Interval Established with Normal Operating Data, Due to Offset of 150 Pa (0.2 in. of water) in Supply Duct Pressure Sensor

nate data points collected during system transients when the supply fan speed was changing to bring the duct static pressure back to the set point.

This method has the desirable property that the threshold for fault detection is determined solely by the selection of a desired confidence interval and the distribution of data during the training phase. For example, unevenly distributed data sets will cause the confidence intervals to expand in airflow regions where data are scarce (Norford and Little 1993). The threshold for fault detection at any flow is therefore not arbitrary but is statistically defensible, under the major assumption that data collected during the training phase represent normal operation of the AHU.

Correlations of power with flow can be used to detect faults involving duct pressure, including pressure sensor faults, stuck dampers, and fouled filters. Figure 2 shows data well outside the confidence intervals, obtained in a fault condition when air was bled from the static pressure

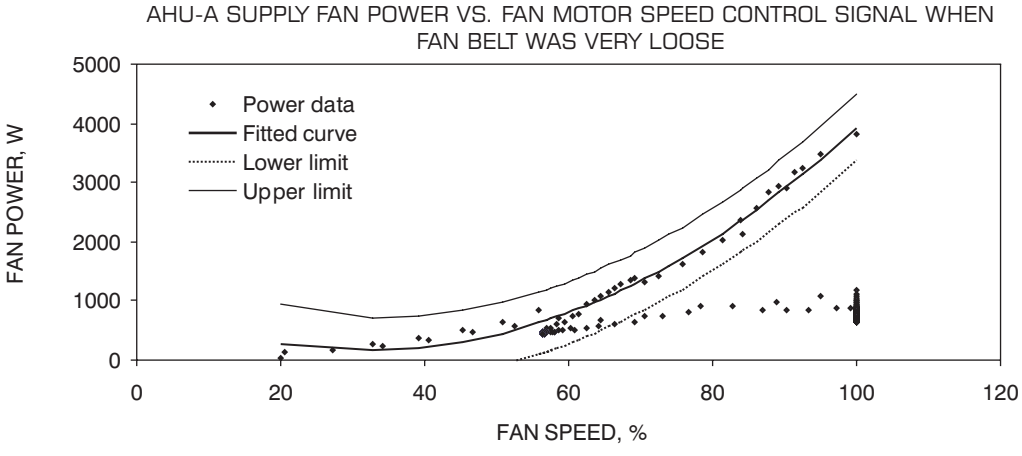

Figure 3. Reduction of Fan Power at Very High Speed-Control Signals, Due to Slipping Fan Belt

Downloaded by [] at 07:55 19 July 2016

sensor's pneumatic line to simulate an offset in the static pressure sensor. The offset was introduced in three stages of 50, 100, and 150 Pa (0.2, 0.4, and 0.6 in. of water), with the most extreme case illustrated in Figure 2.

The correlation shown in Figures 1 and 2 can be used to detect a pressure sensor offset, either positive or negative, at any time of year. The same correlation can also be used to detect the stuck-closed recirculation-damper fault, but only under conditions identified in Table 4, where the damper would normally be wide open or nearly so. Under such conditions, a stuck-closed damper will increase the pressure drop in the air handler and therefore increase fan power for a given airflow.

Correlations of Fan Power with Fan-Motor Speed-Control Signal

Some faults can be detected when power is correlated with motor speed control signal, rather than airflow. For example, reduced tension in a fan belt can cause the belt to slip and transmit less mechanical energy from the motor to the fan at high load. If there is relatively little dissipation of energy in the loose belt, then, for a given motor speed, the motor draws less power and the airflow is lower. Fan motor power is not expected to change significantly for a given airflow. However, the power drawn by the motor for a given rotational speed will be less than normal, because the fan is not spinning as fast as it normally would and therefore transfers less power to the air. This fault can therefore be detected by a deviation from the expected relationship between fan power and the motor speed control signal, as shown in Figure 3. It is important to note that the power-speed correlation is immune to changes in duct pressures that can be detected with the power-flow correlation, as was established in Norford and Little (1993). As a result, the method for detecting the slipping fan belt is also sufficient to distinguish it from the faults detected with the power-flow correlation. However, a power-speed correlation will detect such additional faults as a loose pulley or a transducer failure, making fault diagnosis problematic. Figure 4 shows a transducer failure implemented in the test building.

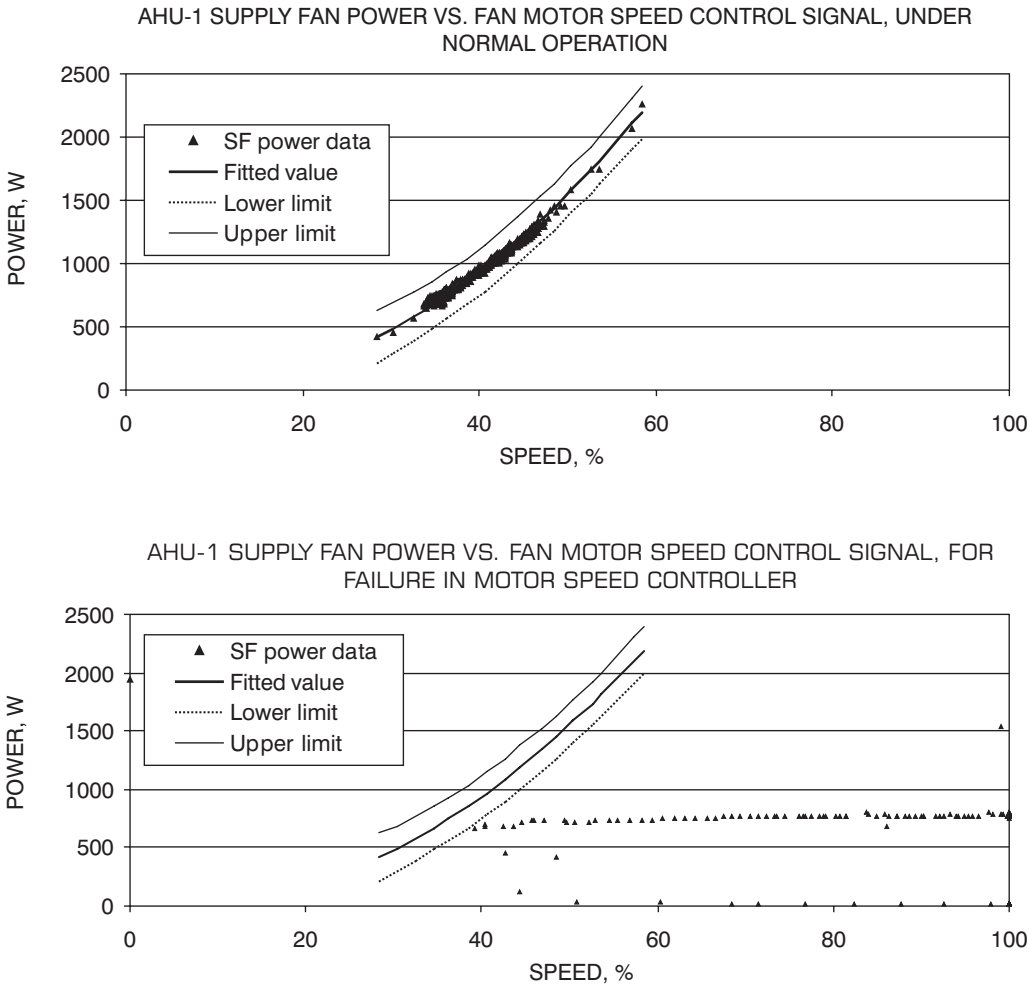

Downloaded by [] at 07:55 19 July 2016

Figure 4. Reduction of Fan Power at High Values of Speed-Control Signal, as Result of Controller Failure

The top figure shows the fit to data from a single day of normal operation. The lower figure shows data during a three-hour period when the fan speed was fixed at a single value, over a range of speed-control inputs.

Correlation of Chilled Water Pump Power with Cooling Coil Valve Position

Changes in pump power can be used to detect blockages in piping, in the same manner as using fan power to detect stuck dampers. In many buildings with a single chilled water loop, the chilled water pump is run at constant speed and flow to the cooling coil is controlled with a three-way valve. The pump will ride the pump curve as flow resistance changes and pump power will change accordingly. Flow resistance can change due to normal operation (operation of the three-way valve to control flow to the coil, which has a larger flow resistance than the by-pass piping unless the latter includes a balancing valve) or due to a fault (flow obstruction). With

a two-way valve in lieu of a three-way valve, pump power will vary more strongly with valve position. While changes in pump power are an attractive means of detecting flow-restriction faults, chilled water flow is rarely measured and therefore is not a desirable correlating variable. The one available control signal, which adjusts the position of the cooling coil valve (either two-way or three-way) was used.

The test building has a primary-secondary piping system, with a variable-speed primary pump and constant-speed secondary pumps for the two test air handlers. Initially, the cooling loops used two-way valves to control flow to the air-handler cooling coils, but these valves were replaced with pairs of two-way valves that are controlled to perform as three-way valves. Further, flow balancing valves were installed to equalize the pressure drop across the coil and the bypass piping. The pump power under normal operation is expected to be nearly independent of valve position. If there is a flow restriction in a cooling coil when balancing valves are installed and used properly, it is expected that there will be less flow through the coil; the discharge air temperature will tend to rise above its set point; the valve controller will open the valve to send more flow through the coil and less to the bypass loop; the overall flow resistance will increase and the pump will ride up the pump curve to a lower total flow and reduced pump power.

Correlating pump power with valve control signal is therefore sufficient to detect flow blockages under cooling loads sufficiently high that a substantial fraction of the total flow is directed through the coil. However, it is not likely to detect coil fouling, where a very thin coating of calcium carbonate can drastically reduce heat transfer across the coil but can have a small impact on flow resistance. The change in pump power as correlated with valve position is illustrated in Figure 5. The training period did not include system loads large enough to cause the cooling coil valve to open more than 70%. Pump power was assumed to remain nearly constant under higher flows through the cooling coil.

Chiller power can provide another indication of faults in chilled water piping. For low to moderate cooling loads, the valve controller will compensate for a flow restriction in the cooling coil by directing more water to the coil and less to the bypass piping. Under high cooling loads, the valve controller will saturate, flow through the coil will be less than needed to maintain the discharge-air temperature, the building will be undercooled, and chiller power will therefore drop. Monitoring to determine a reduction in chiller power at high load also offers the advantage of detecting not only flow-restriction faults but also reduced thermal conductivity due to deposits on the water side of the cooling coil.

Whole-building energy studies have correlated building electricity consumption with outside temperature as a means of analyzing the building's energy requirements for cooling, typically with linear change-point models (Ruch and Claridge 1992, Ruch et al. 1993). More detailed studies of chiller power have established that chiller power is primarily a function of load on the chiller and the temperature difference between leaving condenser water and chilled water flows, and that a biquadratic functional form is a reasonable model (Braun et al. 1987). For an air-cooled chiller as is used in the test building, outside air temperature directly affects condenser performance.

In this study, an HVACSIM+ simulation of a building modeled for a controls simulation test bed (Haves et al. 1996b, 1998) was used to correlate cooling load, as measured by heat transfer across the cooling coil, against outdoor temperature. It was possible to detect a fouled coil. This method was not applied in the test building to detect the coil capacity fault because the chiller is a two-stage reciprocating unit and power levels are discrete, rather than continuously varying. Cycling periods between states were not regular and not easily discerned at high loads. At low cooling loads, however, chiller cycling was both regular and revealing of certain faults, as will be discussed.

Downloaded by [] at 07:55 19 July 2016

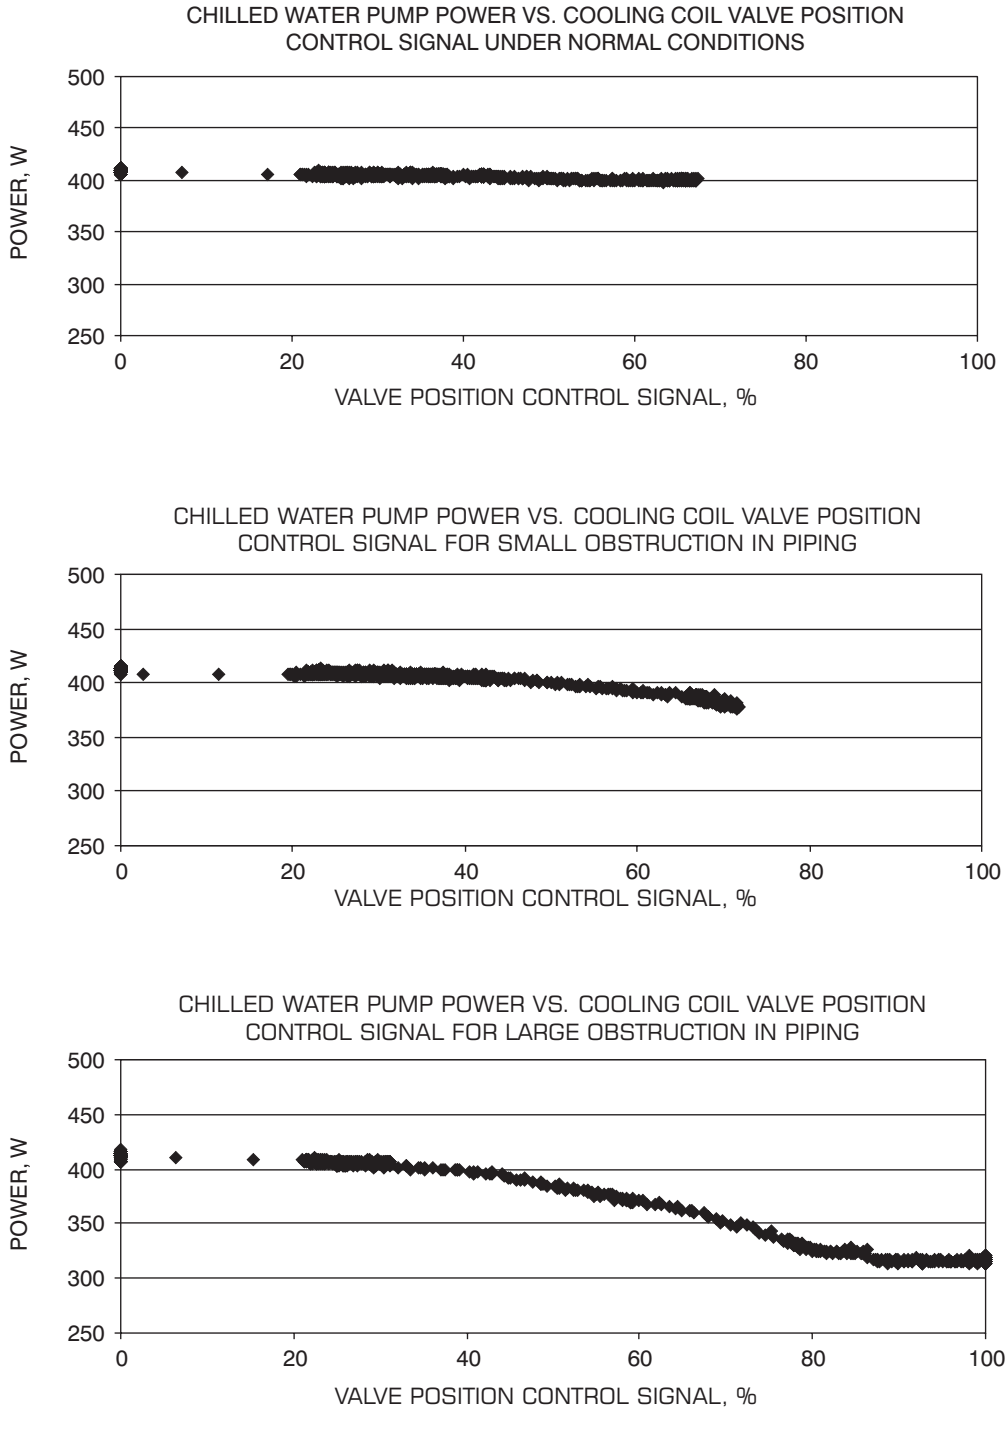

Figure 5. Correlation of Pump Power with Valve Position, for Normal Conditions and for Reduction in Cooling Coil Capacity Due to Obstruction in Chilled Water Piping

Downloaded by [] at 07:55 19 July 2016

Analysis of Reciprocating Chiller Cycling Periods

Cycling periods for the two-stage reciprocating chiller, between shutdown and the low-power state, can be analyzed to detect such faults as a leaky recirculation damper and a leaky cooling coil valve. Both of these faults can be detected only at low cooling loads, when the recirculation damper would normally be shut by the air-side economizer and when the cooling coil valve would be normally closed. Under low loads, a reciprocating chiller will cycle on for a short period and then shut down for a longer period, controlled to maintain the chilled water temperature returning to the chiller within a specified band. For the leaky recirculation damper, the chiller cycling period can be analyzed in a narrow band of outside temperatures, just above the point where the cooling coil valve starts to open and there is a cooling load across the coil. (This corresponds to the boundary between outdoor temperature regions 2 and 3, as defined above.) Limiting the analysis window to a narrow band guards against labeling as a fault those changes in cycling period that are due to a normal increase in cooling load with outside temperature.

This approach works well when outdoor temperatures exceed the supply air temperature set point by a small fraction of the supply-return temperature difference (20% has been used as a cutoff), and, importantly, remain in this region long enough to detect a stable chiller cycling pattern. When the outdoor temperature rises and the building cooling load increases, as is the case on many days, chiller-cycling period is not a good fault detection metric.

The leaky cooling coil valve will increase the load on the chiller for a given valve position. It is difficult to correlate changes in reciprocating chiller cycling periods with valve position, due to the coarse nature of the former. An acceptable alternative is to analyze changes in chiller-cycling period when the valve control signal is zero and the valve is nominally shut. This approach is relatively robust, with less chance for false alarms than if chiller-cycling data were considered when the valve is partially open. It also distinguishes the symptoms of this fault from the leaky recirculation damper.

Figure 6 shows a chiller-cycling period of 20 to 25 min. when the cooling coil valve was closed, much shorter than the normal cycling period of 38 to 39 min. observed earlier the same day and indicative of a leaky valve. The cycling interval dropped to 14 to 15 min. when the cooling coil valve was about 20% open later in the day, a period excluded from analysis.

Downloaded by [] at 07:55 19 July 2016

Analysis of Fan Power Oscillations

Norford and Leeb (1996) showed that centralized power monitoring could detect chiller-power oscillations due to an unstable chilled water temperature controller. Screening for rapid power oscillations forms the basis for detecting faults caused by underdamped or unstable local-loop controllers. Power oscillations are quantified by the standard deviation of the data sampled in a sliding window. In the test building, this approach has been used to detect an unstable pressure-control-loop fault in the test building, as shown in Figure 7.

B. Detection and Diagnosis of HVAC Faults with Gray-Box Models and Centralized Electrical Power Measurements Under Steady-State Conditions

The steady-state, gray-box FDD method was also evaluated when the electrical power data came from centrally located power meters rather than submeters dedicated to individual components. Two centralized meters, known as Non-Intrusive Load Monitors (NILMs), were installed in the test building, one on the motor control center serving five fans and ten pumps, and one at the whole-building electrical service entry. A detailed development of the signal-processing algorithms needed to detect steady-state changes in HVAC loads is presented in Luo et al. (2002). The use of the centralized meters to detect classes of faults, including six of the seven introduced in the test-building air handlers, is briefly described.

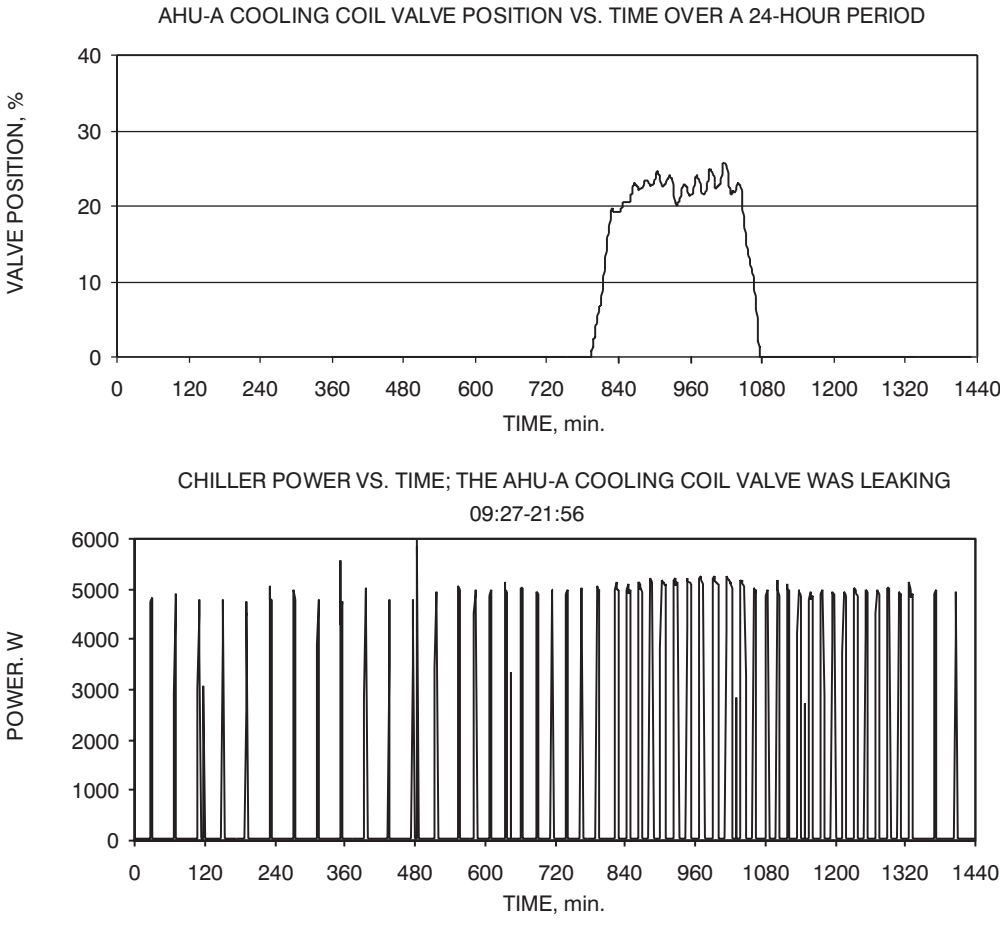

Downloaded by [] at 07:55 19 July 2016

Figure 6. Detection of Leaky Cooling Coil Valve by Analysis of Cycling Period of Reciprocating Chiller

The power-flow correlation for the fan that was readily made with submetered power data could not be generated with electrical power measurements from the NILM installed on the motor control center in the test building. The NILM yielded one data point per day, when the fan was turned off in the evening. The startup point was not valuable because the fan motor has a variable-speed drive that has a slow and complex startup pattern that is very different from an abrupt change most easily seen by the change-of-mean detection algorithm applied to data collected by the NILM. Further, there was little variation in flow at the time of fan shutdown and no opportunity to generate a polynomial relationship between power and flow. The limited range of data also made it impossible to correlate fan power with the motor-speed-control signal, a correlation established with submetered data and used to detect and diagnose the slipping fan belt.

The fact that airflow, fan speed and fan power showed little variation at shutdown from day to day, due to nearly constant (and very light) building loads in late evening, made it possible to use deviation of fan power from the normal value as a basis for fault detection. For example, when the fan belt in the test building slipped severely, the fan power was about 750 W with a

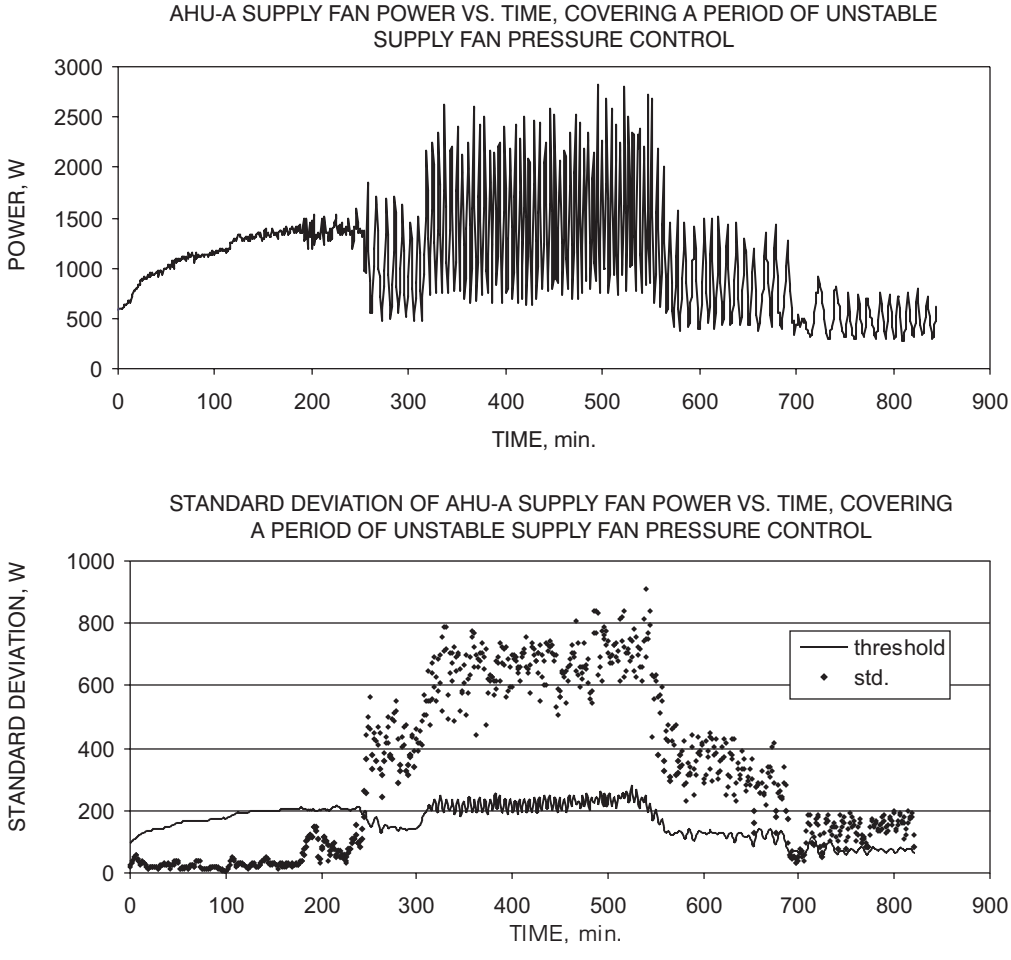

Figure 7. Detection of Unstable Pressure Controller via Analysis of Standard Deviation of Supply Fan Electrical Power

100% motor speed control signal at shutdown, compared to a normal speed of about 20% and a normal power of about 500 W.

Supply fan electrical power of 600 W at shutdown, 100 W above the typical value, was taken as a threshold above which a fault was flagged. This value, selected after examination of the data, proved acceptable in that it did not generate false alarms in the test building. Among the limited number of faults introduced into the test building, the three possible causes were the stuck-closed recirculation damper, an offset in the supply duct static pressure sensor, and the slipping fan belt. Power data alone were sufficient to detect these faults but not to diagnose them. Motor speed data would separate the slipping fan belt from the two faults that are driven by pressure changes.

Figure 8 shows fan power as estimated by the NILM and the supply fan submeter. Power is plotted against airflow to show the limited range of airflows at shutdown, compared with the operating range shown earlier in Figures 1 and 2. The NILM estimate of power is systematically lower than the submeter, a discrepancy that stems from the difference between one-minute

Downloaded by [] at 07:55 19 July 2016

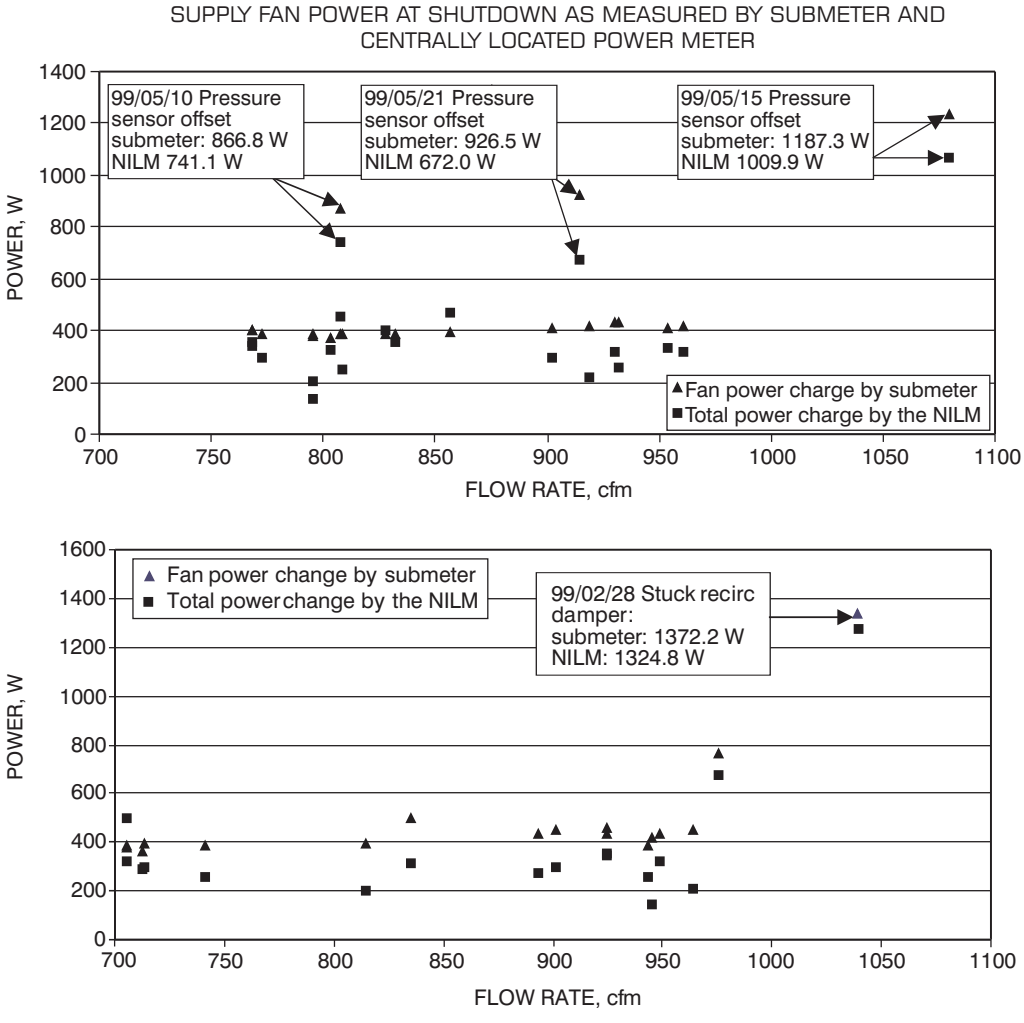

Figure 8. Detection of Pressure Sensor Offset (Top) and Stuck Closed Recirculation Damper (Bottom), via Analysis of Fan Power at Time of Shutdown, as Recorded by Centralized Power Meter with Fault Detection Threshold of 600 W

power averages as used in the submeter and the NILM's calculation of power on the basis of short-term data immediately before and after the shutdown event is detected.

There are two alternative approaches that avoid establishing a somewhat arbitrary power threshold for detecting faults at the time the fan is shutdown. First, fan power as a function of flow and static pressure set point can be modeled from manufacturer's data (Englander and Norford 1992). The model can be tuned with one or more data points. However, such a model does not provide the statistical confidence intervals that come from a polynomial fit of power to flow and a threshold for fault detection would still need to be assigned. Second, fan power data at times other than shutdown can be used to generate the polynomial relation between fan power and airflow or motor speed. Power data can be obtained with a portable power meter used during a commissioning period, or by shutting down the fan several times during a training period,

Downloaded by [] at 07:55 19 July 2016

when airflows or speeds are at values that span the expected range of operations. As with submetered power data, confidence intervals establish a range over which data are considered to be normal and no arbitrary threshold is required. Power at shutdown can then be compared with the expected power for the flow measured at shutdown. This approach was implemented and the power estimate that came from fitting power to airflow during the training period was found to closely predict subsequent power measurements made with submeters.

Submetered power for the secondary chilled water pump was adequate to detect the coil-capacity fault. However, the pump power was so small that it was difficult to reliably detect from the NILM at the motor control center, without lowering the detection threshold to the point where small, unknown, and uninteresting loads were flagged. Further, successful detection of this fault requires accurate analysis of small changes in the small load. Therefore, the NILM was not used to detect pump power.

Luo et al. (2002) describe an effective means for using the meter at the building service entry to pinpoint chiller cycling, even in a very noisy electrical environment. This method boosts sensitivity and reduces false alarms by sampling the power signal at multiple rates, detecting on- and off-events for each data set, and combining the results to reconstruct the sequence of on- and off-events for a given component.

This method has been used to successfully detect both the leaky cooling coil valve and the leaky recirculation damper in the test building. As noted in the discussion of FDD with submeters, the two faults were distinguished by insisting that the cooling coil valve be closed as a prerequisite for diagnosing the leaky valve and by focusing on a narrow range of outdoor temperatures when the recirculation damper should normally be tightly shut and chiller loads are low. Figure 9 shows the use of the NILM at the building service entry to detect the leaky recirculation damper. The middle plot, based on submetered power data, shows a chiller cycling interval of about 15 min. over a 10-h period when the outdoor dry-bulb temperature was sufficiently low. The lower plot shows the output of the NILM. Positive changes in power are associated with the chiller's going from an off-condition to the lower of its two stages, and negative power changes signify shutdown events. The NILM detected all of the transitions, which is what was needed to detect the leaky recirculation damper, although the power the NILM associated with these transitions differs from the submetered value of 5 kW. The normal cycling period was observed to be well above the 30-min. threshold established for this fault.

Power oscillations were detected with the NILM attached to the motor control center, indicative of the unstable-supply-duct-pressure-controller fault introduced in the test building. In general, fault diagnosis from a central point is not possible. While it may be possible to assign the magnitude of the power oscillations to classes of equipment (fans, pumps, chillers), it is not possible to pinpoint a particular fan or pump if there are multiple units of comparable size.

C. Detection and Diagnosis of HVAC Faults with Dynamic Models and Submetered Electrical Power Measurements

The chilled water pump and the supply fan associated with AHU-B in the test building were considered for detection and potential diagnosis of faults using a dynamic-modeling approach. The method used involves fitting parameters to a physical model of the startup behavior of the motor and its attached load using only current measurements available at the stator of the motor. Changes in these physical parameters provide useful HVAC diagnostic information such as flow obstructions.

Pump Diagnostics

Deposits in the pipes of a heater exchanger are difficult to detect noninvasively and can contribute to decreased heating and cooling efficiency. In principle, it might be possible to detect

Downloaded by [] at 07:55 19 July 2016

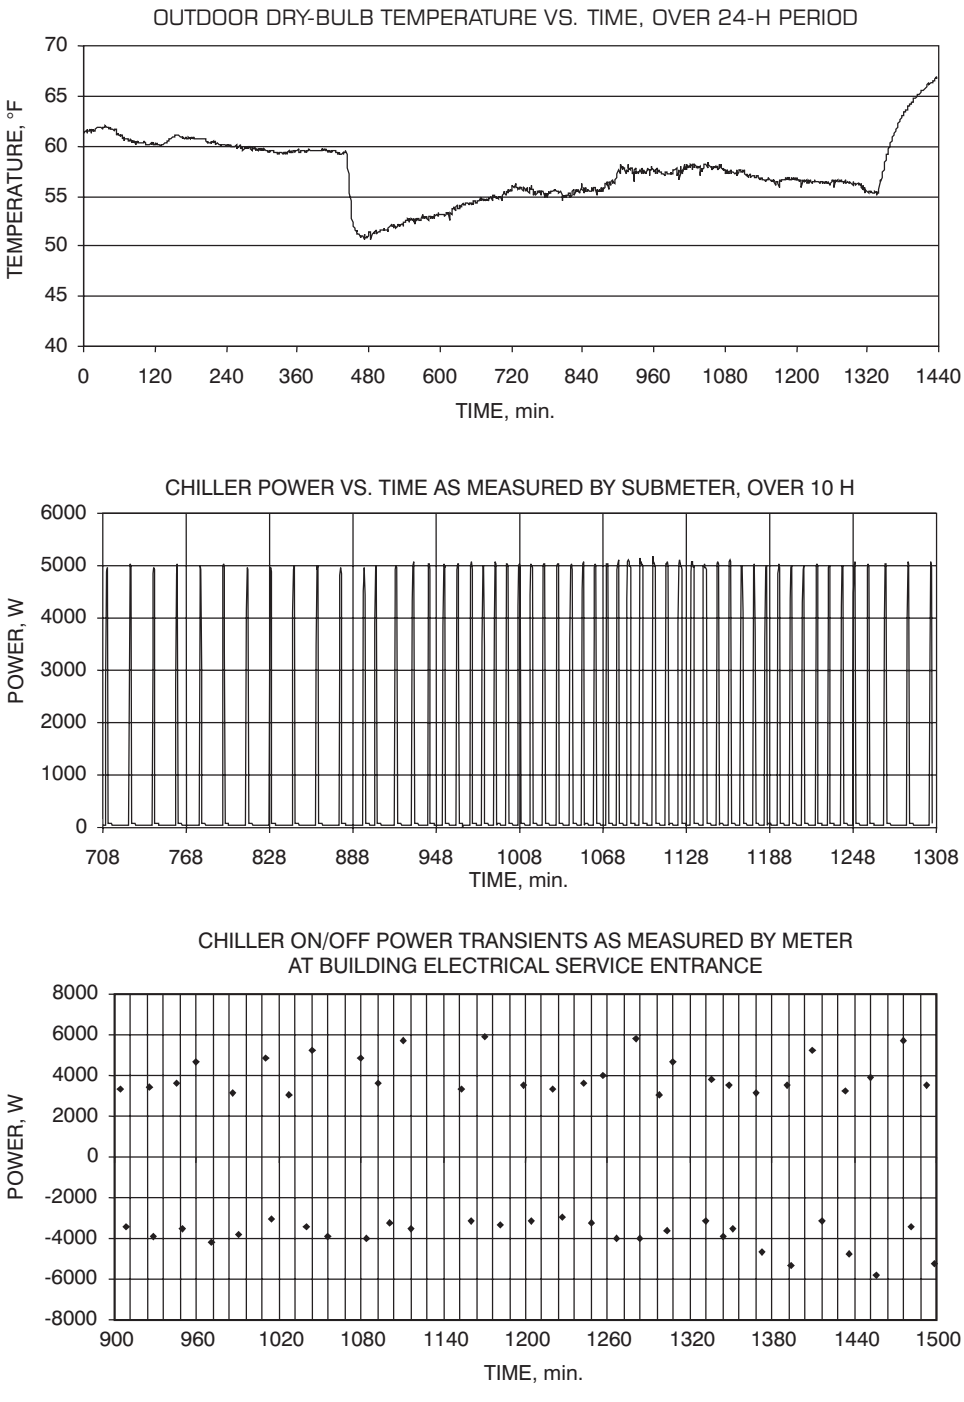

Figure 9. Detection of Leaky Recirculation Damper via Analysis of Reciprocating Chiller Cycling Intervals with Outdoor Temperature from 13 to 15°C (55 to 59°F)

The second plot shows submetered chiller power and the lower plot reveals the power changes (positive for turn-on, negative for turn-off) detected by the NILM.

Downloaded by [] at 07:55 19 July 2016

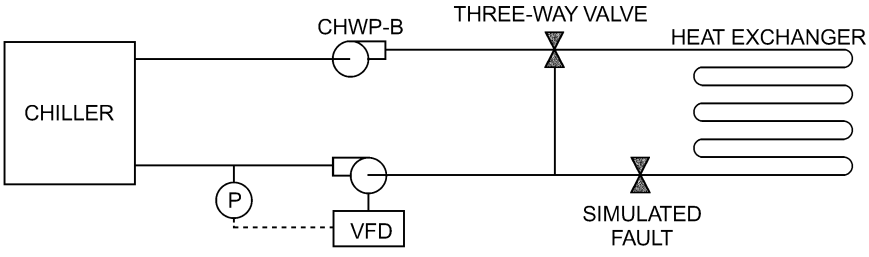

Figure 10. Simplified Diagram of Cooling Loop for AHU-B in Test Building
A single variable-speed primary pump and a fixed-speed secondary pump were used for each air handler. An obstructed cooling coil was simulated with a throttling valve. Current to the secondary pump CHWP-B was measured.

deposits by estimating flow resistance for the fluid in the pipes. Electrical transients were recorded for a chilled water pump in the HVAC system for both normal and an obstructed-flow condition. Detection of the obstructed flow is demonstrated via parameter estimation using the collected data.

A simplified diagram of the chilled water circulation system is shown in Figure 10. The chilled water system consists of a single, variable-speed primary pump and one fixed-speed secondary pump for each air handler. For the test, the three-way cooling coil valve was positioned so that all liquid flowed through the heat exchanger. The primary chilled water pump, equipped with a variable-frequency drive, was operated to control pressure as indicated. The response of this control loop was presumed to be slow enough to ignore in comparison to the startup transients of the secondary pump, CHWP-B. Further, the impact of the primary loop on the secondary pump startup transient could be accurately modeled by including the fixed pressure of the primary loop and not the complete primary loop flow path. To simulate the obstructed flow fault, the building operators installed a valve in series with the heat exchanger, as indicated. In the no-fault condition, with both primary and secondary pumps running, flow in the loop was approximately 1.7 L/s (27 gpm). With the fault in place, flow was reduced to 0.7 L/s (11 gpm), which is approximately 40% of nominal flow.

The experimental procedure was to introduce the simulated fault, turn CHWP-B on and off a few times, remove the fault, and again cycle CHWP-B. The resulting startup transients were recorded and then analyzed off-line.

Transducers were installed to measure current on two of the phases feeding the balanced, three-phase 480 V pump motor. Data were sampled with 12-bit resolution at 4000 Hz. A typical transient is shown in Figure 11. Inspection of the data revealed the presence of an unanticipated fault, consisting of an occasional line-cycle of zero current in the transient. This phenomenon was attributed to the contactor. Because contactor failure was not part of the model, transients with contactor problems were discarded. The remaining transients were labeled A, B, C, and D and will be referred to as such when results are presented. Transients A and B were collected under fault conditions.

A simple electrical model for the cooling system is an induction machine connected to an inertia with damping. In addition to the six parameters implied by this simplistic model, it is necessary to estimate the electrical angle when the motor was turned on and the initial speed of the motor. Extracting eight parameters from a simple transient as shown in Figure 11 is difficult.

It was assumed that the fault lay in the mechanical system (i.e., that the motor performance was the same when comparing two transients). A “joint,” two-transient model was formed

Downloaded by [] at 07:55 19 July 2016

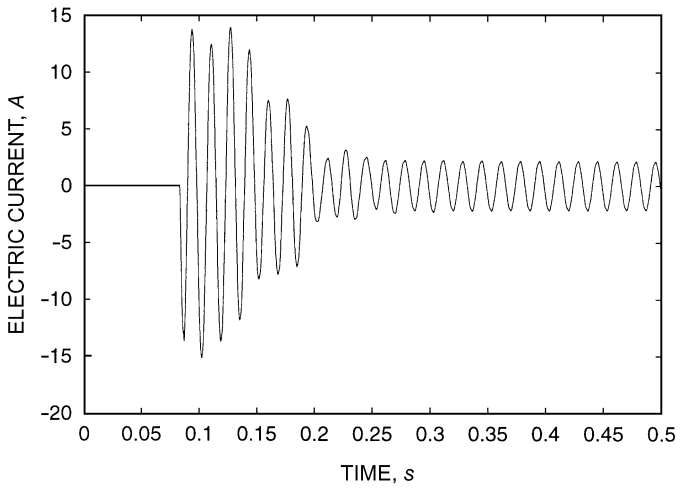

Figure 11. Typical Electrical Transient for CHWP-B Motor
The transient is the fault transient labeled "A" in Table 5.

Downloaded by [] at 07:55 19 July 2016

where identical motor parameters were used for both transients, while the mechanical parameters were allowed to differ for each transient. Mathematically, the model consists of induction motor equations in what is known as *dq* space:

$$\frac{d}{dt}\begin{pmatrix}\lambda_{qs}\\ \lambda_{ds}\\ \lambda_{qr}\\ \lambda_{dr}\end{pmatrix} = \begin{pmatrix}\nu_{qs}\\ \nu_{ds}\\ 0\\ 0\end{pmatrix} - \begin{pmatrix}r_s i_{qs} + \omega\lambda_{ds}\\ r_s i_{ds} - \omega\lambda_{qs}\\ r_r i_{qr} + (\omega - \omega_r)\lambda_{dr}\\ r_r i_{dr} - (\omega - \omega_r)\lambda_{qr}\end{pmatrix} \tag{1}$$

where ω is the electrical frequency, ω_r is the rotor frequency, and the λs are the flux linkages (Krause 1986). The *dq* transform is a change of coordinates to a rotating frame, often applied to variables in electric machines to simplify analysis. The flux linkages are related to the currents by

$$\begin{aligned}\lambda_{qs} &= L_l i_{qs} + (L_l + L_m)(i_{qs} + i_{qr})\\ \lambda_{ds} &= L_l i_{ds} + (L_l + L_m)(i_{ds} + i_{dr})\\ \lambda_{qr} &= L_l i_{qr} + (L_l + L_m)(i_{qs} + i_{qr})\\ \lambda_{dr} &= L_l i_{dr} + (L_l + L_m)(i_{ds} + i_{dr})\end{aligned} \tag{2}$$

These electrical dynamics are influenced by a set of mechanical interactions, one for each transient:

$$\begin{aligned}\frac{d}{dt}\omega_1 &= 3K_1(\tau_1 - \beta_1\omega_1)\\ \frac{d}{dt}\omega_2 &= 3K_2(\tau_2 - \beta_2\omega_2)\end{aligned} \tag{3}$$

where the torque of electrical origin for the subscripted system is given by

$$\tau = \lambda_{qr} i_{dr} - \lambda_{dr} i_{qr} \tag{4}$$

The term β_1 (β_2) is a damping parameter and K_1 (K_2) is a parameter inversely proportional to the rotational inertia of the motor and pump. The output equations bring the simulation variables to the measured currents in the lab frame of reference. For each transient,

$$i = i_{ds}\cos[\omega(0)t + \phi] - i_{qs}\sin[\omega(0)t + \phi] \tag{5}$$

where ϕ is a parameter representing the electrical turn-on phase for each transient. The complete parameter vector of each joint model is then

$$\mu = [r_r r_s L_m L_l K_1 \beta_1 \omega(0)_1 \phi_1 K_2 \beta_2 \omega(0)_2 \phi_2]' \tag{6}$$

This parameter vector is then optimized to fit both transients in a given data set, as listed in Table 5. The result is that different mechanical parameters are obtained for each transient, while maintaining common electrical parameters. This model structure encapsulates the assumption that the difference between the two transients lies in the mechanical system. Parameter identification with different combinations of transients provides a useful crosscheck of the consistency of the estimates.

The model fit the data well, as shown in Figure 12. The residuals were small, but not structureless, as revealed by such typical residual analysis tools as the zero-crossing test and the Kolmogorov-Smirnov test (Johansson 1993). The model is simple and does not include complicated modeling of the interaction between water and pump. The residuals indicate that a more complicated model might be feasible with the data available. On the other hand, a more complicated model might cause identification problems with different data sets.

Parameters for the six data sets are shown in Tables 6 and 7. The data sets are separated by type, as given in Table 5. The parameters in Tables 6 and 7 agree well. In particular, the individual transients produce remarkably consistent estimates independent of the transients that they are paired with when fitting a joint model.

It is not necessarily important that the residuals are not structureless. The key issue for diagnostic purposes is the robustness of the parameters under perturbations of the measurements (i.e., whether parameter values are interpreted as faulty or not depends on the disturbances). This issue can be addressed by making parameter distribution estimates using synthetic data sets. Synthetic data set estimates of the distributions of the mechanical parameters, under fault and no-fault conditions with data set P_0, appear in Figure 13, which strongly suggests that detection of the fault would be successful under the presumed disturbance.

Table 5. Data Set Organization for Pump Tests

Data Set	Transient 1	Transient 2	Type
P_0	A	C	Fault/No-fault
P_1	A	D	Fault/No-fault
P_2	B	C	Fault/No-fault
P_3	A	D	Fault/No-fault
P_4	A	B	Fault/Fault
P_5	C	D	No-fault/No-fault

Downloaded by [] at 07:55 19 July 2016

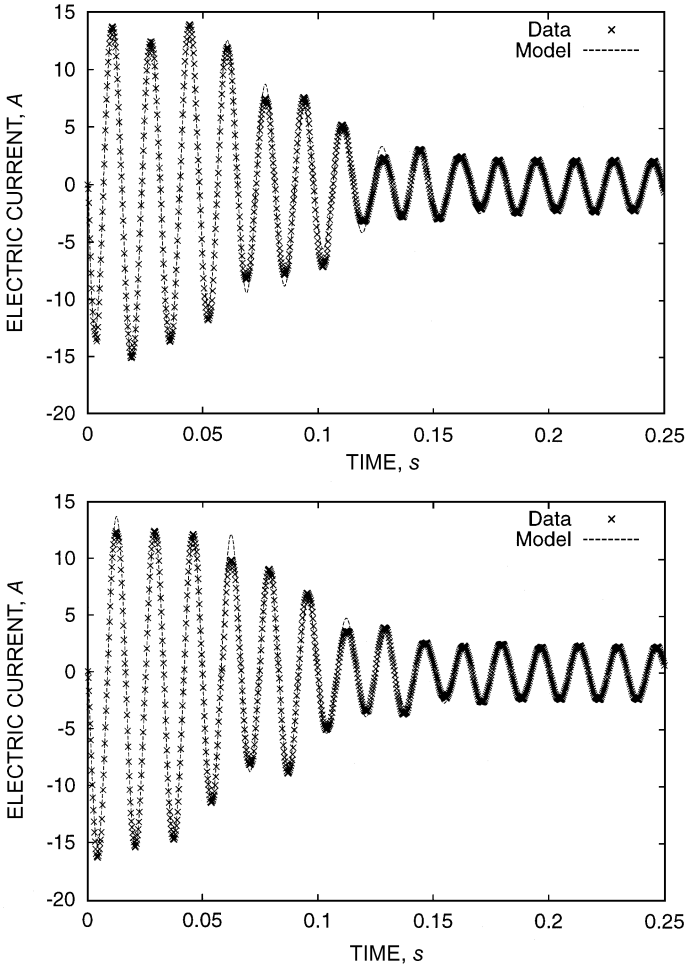

Figure 12. Comparison of Model to Experimental Data for Data Set P_0

Data with the reduced-flow fault are shown in the upper figure, while fault-free data are in the lower figure. The quality of fit is typical of the other data sets. While the fault-free and fault data appear visually similar, the differences are sufficient to produce significantly different values for the mechanical parameters.

Table 6. Electrical Parameter Estimates by Data Set for Pump Tests

	Transient		Electrical Parameters			
Data Set	**1**	**2**	r_r	r_s	L_m	L_l
P_0	A	C	1.548e + 01	1.424e + 01	7.266e – 01	3.046e – 02
P_1	A	D	1.487e + 01	1.418e + 01	7.050e – 01	3.096e – 02
P_2	B	C	1.529e + 01	1.421e + 01	7.205e – 01	3.061e – 02
P_3	A	D	1.511e + 01	1.418e + 01	7.132e – 01	3.077e – 02
P_4	A	B	1.545e + 01	1.380e + 01	7.324e – 01	3.072e – 02
P_5	C	D	1.499e + 01	1.462e + 01	7.048e – 01	3.046e – 02

Downloaded by [] at 07:55 19 July 2016

Downloaded by [] at 07:55 19 July 2016

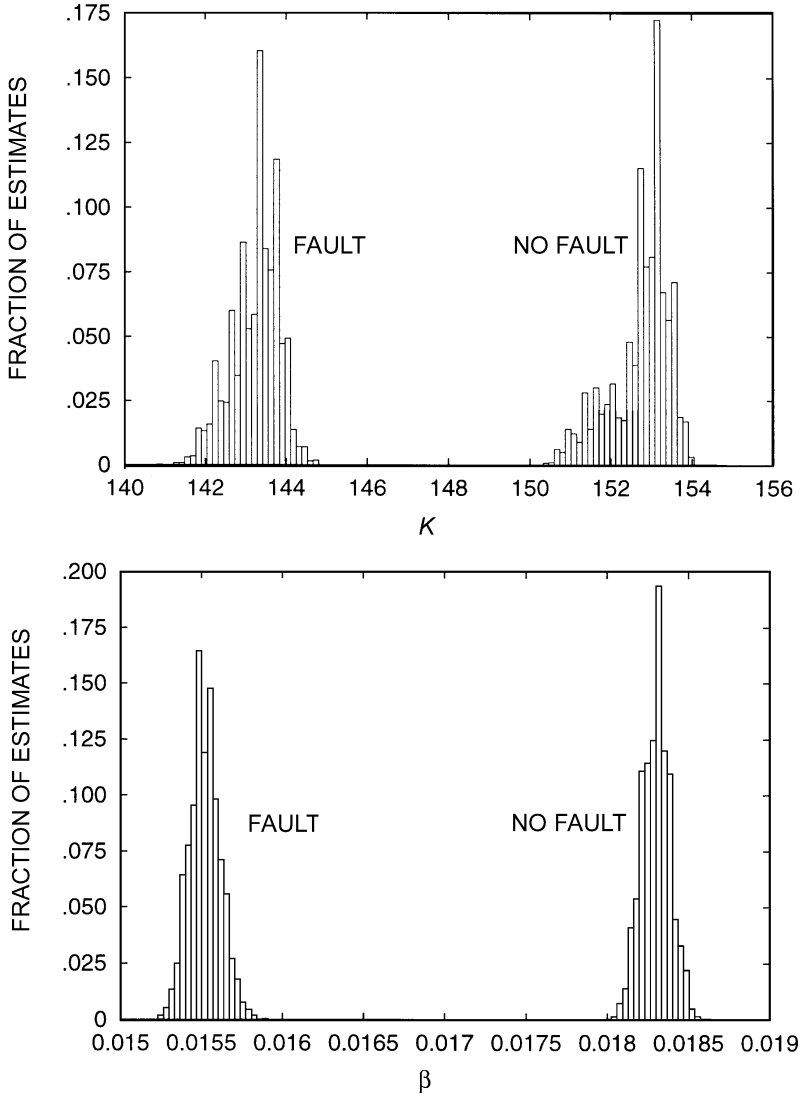

Figure 13. Distribution Estimates for Mechanical Parameters β and *K*

Table 7. Mechanical Parameter Estimates by Data Set for Pump Tests

	Transient		Mechanical Parameters			
Data Set	**1**	**2**	K_1	β_1	K_2	β_2
P_0	A	C	1.435e + 02	1.544e – 02	1.531e + 02	1.830e – 02
P_1	A	D	1.415e + 02	1.498e – 02	1.516e + 02	1.771e – 02
P_2	B	C	1.427e + 02	1.515e – 02	1.531e + 02	1.828e – 02
P_3	A	D	1.422e + 02	1.534e – 02	1.524e + 02	1.783e – 02
P_4	A	B	1.415e + 02	1.566e – 02	1.454e + 02	1.540e – 02
P_5	C	D	1.496e + 02	1.803e – 02	1.487e + 02	1.763e – 02

Downloaded by [] at 07:55 19 July 2016

Fan Diagnostics

To investigate diagnostic possibilities in ventilation systems, data were collected from one of the test air-handling units. A diagram of the air-handling unit appears in Figure 14. Transients from the motor connected to supply fan B were measured, and two kinds of faults were considered. Unlike the pump situation, two functioning current transducers were available for measurement. For the first diagnostic, the outdoor air damper (indicated on the lower right of Figure 14) was opened to create a gross change in the flow characteristics of the system. The second diagnostic concern in the air-handling unit was the effect of a slipping fan belt, introduced in the test building by moving the motor to loosen the belt.

Developing an identification model of a complicated situation like Figure 14 is a challenging task. As a first assumption, controls (e.g., for the return fan) were assumed to operate on a time scale much slower than the induction motor startup transient. Also, other aspects of the system (e.g., state of doors and windows in rooms served by the air-handling unit) were assumed to remain constant during the test. The joint modeling technique used in the pump diagnostics was also used (i.e., the induction motor was assumed to be a constant for purposes of comparing two situations). The mechanical situation was modeled as in the pump scenario, i.e.,

$$\begin{aligned} \frac{d}{dt}\omega_1 &= 3K_1(\tau_1 - \beta_1\omega_1) \\ \frac{d}{dt}\omega_2 &= 3K_2(\tau_2 - \beta_2\omega_2) \end{aligned} \tag{7}$$

Unlike the pump diagnostics, two phases of current measurements were available for fan transients. These measurements were transformed to synchronously rotating *dq* space (Krause 1986) for identification.

Parameter estimates for the fan tests are given in Tables 8 and 9. As with the pump tests, transients were analyzed in pairs to detect changes in mechanical parameters associated with the presence of faults. Startup transient A refers to the no-fault situation, B corresponds to the open outdoor air damper (simulated by opening the mixing box door), C is the slipping belt, and D is both slipping belt and open door. The electrical parameters of Table 8 should be essentially the

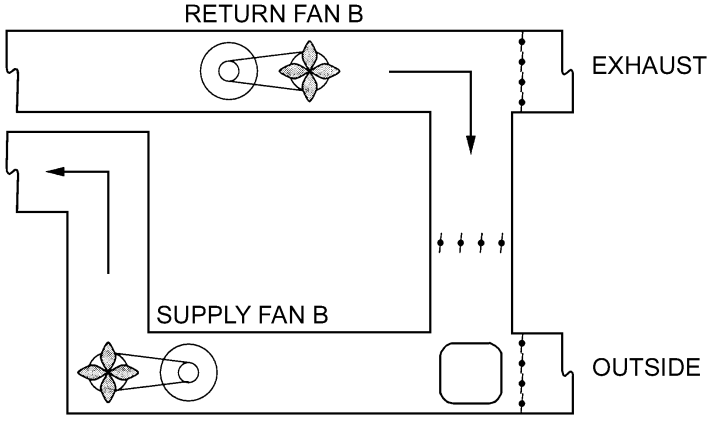

Figure 14. Schematic Diagram of AHU-B

Arrows indicate direction of airflow. Dampers on the right allow building air to be recirculated or mixed to varying degrees with outside air. The supply fan was used in the tests.

same for all combinations of data sets, because the same motor was used throughout. There is quite close agreement for parameters r_s and L_l and not such good agreement for the magnetizing inductance L_m and rotor resistance r_r. It is possible that the rotor resistance r_r might vary due to thermal effects, but the variation seen in the magnetizing inductance L_m is extreme. It seems likely that r_r and L_m are being influenced by the changes introduced in the mechanical part of the system. The most suspect combination in Table 8 is the first row.

The mechanical parameters for the six combinations of data sets are given in Table 9. In this table, parameters with subscript *i* correspond to the *i*th transient column. For example, β_2 in row three corresponds to an estimate for data set D (open door, loose belt) in conjunction with data set A (no-fault). Because the transients appear in different combinations in Table 9, a cross-validation check of sorts can be performed. Mechanical parameters corresponding to an individual transient (A, B, C or D) should be roughly the same. Figure 15 aids in this comparison by plotting the parameters corresponding to different mechanical situations. Note that the "no-fault" points in Figure 15 are quite distinct from the "fault" points. Also, the outlier in Figure 15 corresponds to the suspect row F_0 of Tables 8 and 9.

The parameters in Figure 15 and Table 9 make good physical sense. The situations where the fan belt is slipping show the lowest friction coefficient. As shown in Figure 14, the slipping fan belt situation is relatively close to the slipping fan belt and open door scenario. If the motor is not coupled to the fan, the duct configuration is irrelevant. Physical interpretation of *K* in Figure 15 is more involved. Because *K* is inversely proportional to the rotational inertia of the motor-fan system, larger values of *K* mean *less* inertia. The open door scenario has the most coupled inertia, followed by both slipping fan belt cases, and the no-fault situation has the least inertia. Because the slipping fan belt essentially uncouples the motor (at least for the startup transient), the motor should see only the inertia of the rotor and pulley. This hypothesis agrees well with

Downloaded by [] at 07:55 19 July 2016

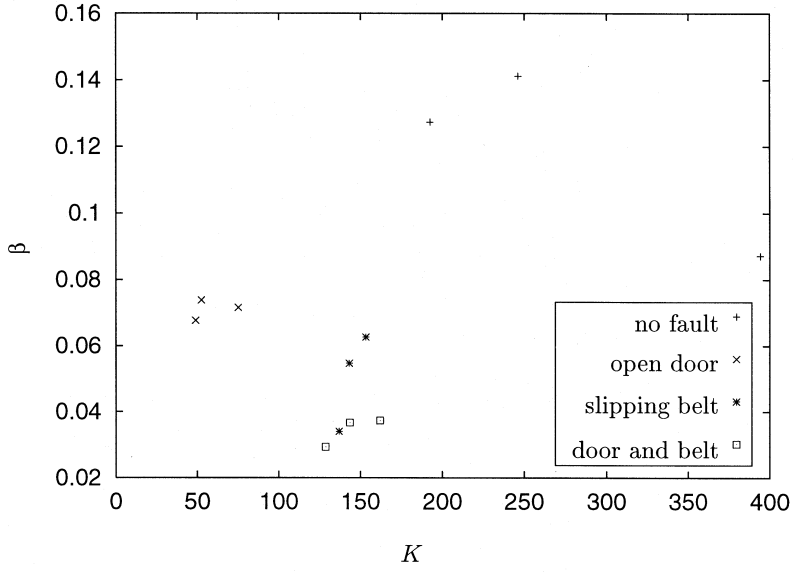

Figure 15. Scatter Plot Showing Parameters for Each Data Set When Estimated in Conjunction with Other Data Sets

The nine points in the lower left corner are fault points. Note that the rightmost point corresponds to the combination in the first row of Tables 8 and 9, which is suspect.

Table 8. Electrical Parameter Estimates for Fan Tests

	Transient		Electrical Parameters			
Data Set	1	2	r_r	r_s	L_m	L_l
F_0	A	B	2.35e – 01	8.07e – 01	5.84e – 02	4.62e – 03
F_1	A	C	3.75e – 01	8.15e – 01	9.72e – 02	5.13e – 03
F_2	A	D	3.99e – 01	7.95e – 01	9.14e – 02	4.99e – 03
F_3	B	D	2.89e – 01	7.98e – 01	1.17e – 01	4.97e – 03
F_4	C	B	3.02e – 01	8.06e – 01	9.79e – 02	4.58e – 03
F_5	C	D	3.26e – 01	8.03e – 01	1.52e – 01	4.78e – 03

Table 9. Mechanical Parameter Estimates for Fan Tests

	Transient		Mechanical Parameters			
Data Set	1	2	K_1	β_1	K_2	β_2
F_0	A	B	3.94e + 02	8.72e – 02	7.51e + 01	7.16e – 02
F_1	A	C	2.46e + 02	1.41e – 01	1.37e + 02	3.40e – 02
F_2	A	D	1.92e + 02	1.27e – 01	1.28e + 02	2.93e – 02
F_3	B	D	5.25e + 02	7.38e – 02	1.62e + 02	3.73e – 02
F_4	C	B	1.53e + 02	6.26e – 02	4.89e + 01	6.76e – 02
F_5	C	D	1.43e + 02	5.46e – 02	1.43e + 02	3.67e – 02

Downloaded by [] at 07:55 19 July 2016

the close inertia values obtained for both slipping belt situations. In the no-fault scenario, the inertia values are *lower* than in the uncoupled case. This makes sense because the no-fault case is the only case where the return fan is supplying a force that tends to accelerate the rotor. The back pressure should not translate exactly into a change in inertia, but its gross effect may, particularly in the absence of a more complex model. When the outdoor air damper is open, the coupling of the supply fan to the return fan is greatly lessened, and the effective inertia includes the air column coupled to the fan.

The fit of the model to the observations was fairly good (see Figure 16). Note that these graphs show data that have been *dq* transformed (Krause 1986). Other mechanical models were attempted, including adding a $\gamma\omega^2$ term to Equation (5). These models did not offer much improvement in the fit, and in some cases it was clear that the data would not support a more complicated model.

CONCLUSION

Two methods have been developed for using electrical power measurements to detect a number of HVAC faults. The first, gray-box, method correlates steady-state power with such exogenous variables as flow, motor speed, or actuator control signal. Data that fall outside a tolerance band around a polynomial correlation indicate a fault. Faults can be diagnosed in limited cases where there are no other likely faults with the same power signature. The first method also detects and diagnoses cooling load faults via analysis of the cycling time of a reciprocating compressor. The second method fits high-speed power measurements taken during motor startup transients to a physical model that includes both electrical and mechanical parameters. Changes in the mechanical parameters indicate a fault. Both methods were demonstrated with data taken at a test building, where a modest number of artificial faults were introduced in the air handlers.

The gray-box FDD method was implemented with data from submeters and also, to a lesser extent, with data from centralized, high-speed electrical load monitors. Fault detection with this method has been limited to a relatively small range of possible faults in an air handler. Some

Downloaded by [] at 07:55 19 July 2016

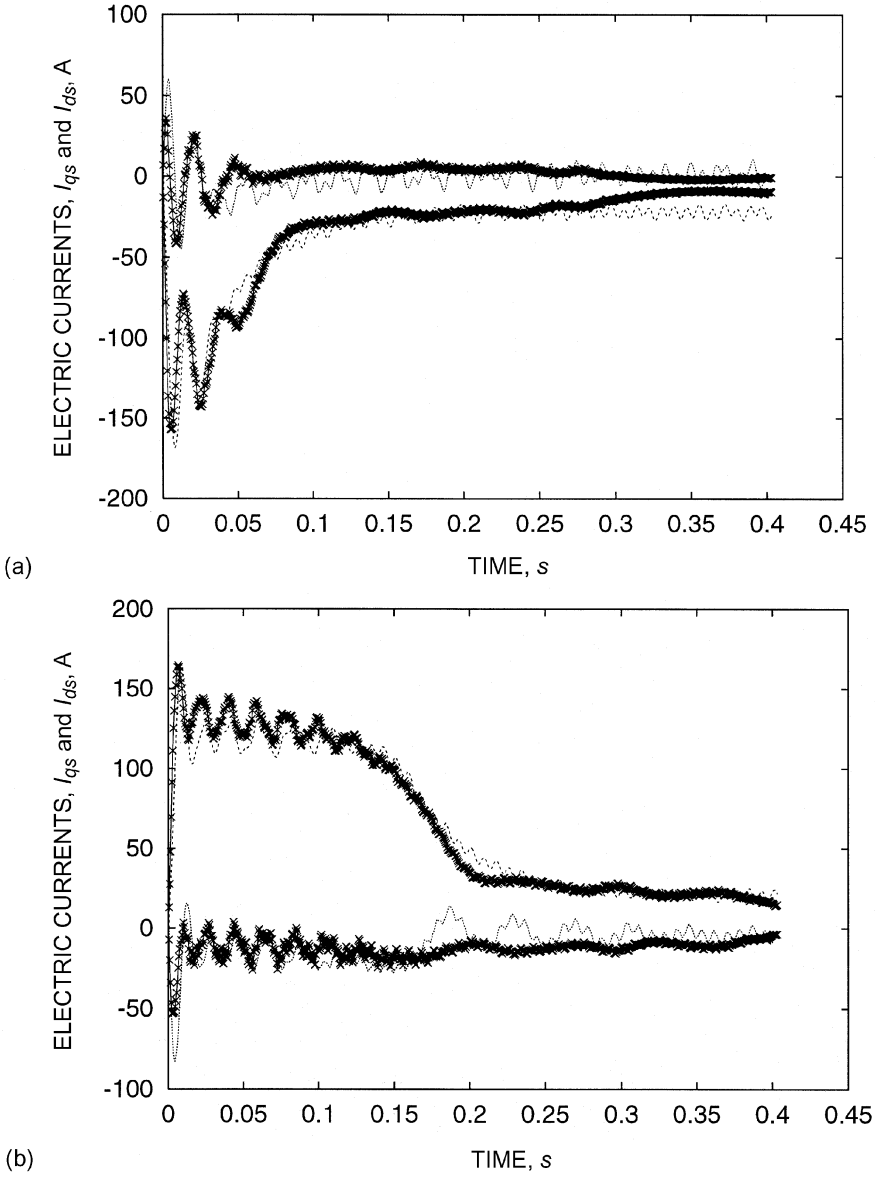

Figure 16. Fit of Joint Model to Experimental Data Sets
(a) Tight belt and 100% recirculation and (b) Tight belt and open outdoor air damper

forms of power correlations show more scatter than others and are less useful. For example, a correlation of fan power with flow has proved to be helpful in detecting changes in flow resistance, but it is more difficult to tightly correlate chiller power with outside temperature as a means of detecting a fault in an outdoor air temperature sensor. The strength of the correlation will have a strong impact on the extent to which the method can be extended to detect and diagnose a larger set of faults.

Detection and especially diagnosis capabilities were reduced when using a centralized, high-speed power meter. The centralized meter unavoidably has less resolution than dedicated sub-

meters, especially for relatively small pieces of equipment. Further, establishing the power correlations proved to be difficult. The central meter only collects data at equipment startup and shutdown. At the test building used in this study, startup signals for AHU fans were masked by the very slow control loop within the variable-speed drives. The secondary chilled water pump ran continuously, producing no power transients. Further, airflow, fan speed, and fan power at shutdown showed little deviation from day to day, due to similar system loads in late evening when the air handlers were turned off. Therefore, the central meter was used only to compare estimated fan power against a value considered to be normal, on the basis of shutdown measurements over a training period. Using only the power measurement made for an efficient means of fault detection but diagnosing the cause of the power deviation was not possible. For example, a deviation in fan power at the normal end-of-day shutdown could be due to a stuck damper or pressure sensor error. An alternative approach, now under investigation, requires that the fans be shut down during the day under different loads. With power data from the central meter and corresponding flow or speed data, power correlations can be established in the same manner as was done with submetered data.

While the use of electricity data for FDD was initially motivated by reduction in metering costs associated with the central meter, the first FDD method may be more appropriately implemented with submetered data if fault diagnosis and not just fault detection is highly valued. Less expensive sensors and sensor-communication systems (e.g., wireless sensors, to avoid wiring costs) will facilitate the use of submetered data.

Downloaded by [] at 07:55 19 July 2016

The second method relies solely on submetered data and a physical model, which is more difficult to formulate than a simple power correlation but is also more powerful. It needs no correlative data, such as flow measurements, and can detect a fault with data collected over the very short period of time required for a motor to reach a steady speed. However, faults can only be detected at the time of motor startup, which introduces an unwanted lag between fault introduction and fault detection. Extending the method to work with data from a centrally located electrical meter, as will be attempted in the future, is possible in principle if the startup transient for a single motor is not masked by electrical noise introduced by other components. Limited laboratory tests have shown that a startup transient can be properly associated with a given motor even when another motor's startup partially overlaps the transient under investigation. However, such tests have not accounted for the sheer volume of events in a real building. Further, variable-speed drives are potentially a major stumbling block, because their prolonged startup cycle further masks the dynamics of the motor and driven load. This method is well suited to submetered power data and may appeal in the future as an onboard diagnostic method for equipment provided with a low-cost current transducer.

Both methods are potentially more robust than FDD methods that rely on temperature and flow sensors, in the sense that they do not require estimating small temperature differences with sensors that are subject to errors. Power measurements also avoid concerns about placement of flow and temperature sensors in large ducts or pipes. Both methods require testing in more buildings.

ACKNOWLEDGMENT

The authors warmly acknowledge the financial support of Honeywell, Inc., the California Energy Commission (via subcontracts from Lawrence Berkeley National Laboratory and Architectural Energy Corporation), and ASHRAE; the vision and guidance of Phil Haves, John House, and Kristin Heinemeier; and the expertise and dedication of Curt Klaassen and Andy Suby at the Iowa Energy Center's Energy Resource Station, the test building used in our field work.

NOMENCLATURE

T_{out}	outside dry-bulb temperature
$T_{balance\ point}$	balance point temperature, the outside dry-bulb temperature at which a building requires neither heating nor cooling
$T_{supply\ air}$	temperature of the supply air
$\Delta T_{supply\ fan}$	temperature rise across the supply fan
$T_{return\ air}$	temperature of the return air
λ	flux linkage for an induction motor
ω	frequency of excitation at the stator of an induction motor
ω_r	frequency of the rotor of an induction motor
i_{dr}	direct rotor current in the *dq* frame
i_{ds}	direct stator current in the *dq* frame
i_{qr}	quadrature rotor current in the *dq* frame
i_{qs}	quadrature stator current in the *dq* frame
L_i	leakage inductance of an induction motor
L_m	magnetizing inductance of an induction motor
r_r	electrical resistance of the rotor of an induction motor
r_s	electrical resistance of the stator of an induction motor
υ_{ds}	direct stator voltage in *dq* frame
υ_{qs}	quadrature stator voltage in *dq* frame
β	damping parameter for a system consisting of a motor and pump or fan
K	parameter inversely proportional to the rotational inertia of the motor and pump or fan
ϕ	the electrical turn-on phase for a motor startup transient
NILM	non-intrusive electrical load monitor

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* 120:224-229.

Ahn, B.-C., J.W. Mitchell, and I.B.D. McIntosh. 2001. Model-Based Fault Detection and Diagnostics for Cooling Towers. *ASHRAE Transactions* 107(1).

Braun, J.E., J.W. Mitchell, S.A. Klein, and W.A. Beckman. 1987. Performance and Control Characteristics of a Large Cooling System. *ASHRAE Transactions* 93(1).

Chen, B. and J.E. Braun. 2001. Simple Rule-Based Methods for Fault Detection and Diagnostics Applied to Packaged Air-Conditioners. *ASHRAE Transactions* 107(1).

Dexter, A.L. and M. Benouarets. 1996. A Generic Approach to Identifying Faults in HVAC Plan. *ASHRAE Transactions* 102(1):550-556.

Dexter, A.L. and D. Ngo. 2001. Fault Diagnosis in Air-Conditioning Systems: A Multi-Step Fuzzy Model-Based Approach. *International Journal of HVAC&R Research* 7(1):83-102.

Englander, S. and L. Norford. 1992. Variable Speed Drives: Improving Energy Consumption Modeling and Savings Analysis Techniques. *Proceedings of the 1992 ACEEE Summer Study on Energy Efficiency in Buildings,* Pacific Grove, California 3:61-78.

Fels, M.F. 1986. PRISM: An Introduction. *Energy and Buildings* 9(1,2):5-18.

Hart, G.W. 1992. Nonintrusive Appliance Load Monitoring. *Proceedings of the IEEE* 80(12):1870-1891.

Haves, P., T.I. Salsbury, and J.A. Wright. 1996a. Condition Monitoring in HVAC Subsystems Using First-Principles Models. *ASHRAE Transactions* 102(1):519-527.

Haves, P., L.K. Norford, M. DeSimone, and L. Mei. 1996b. *A Standard Simulation Testbed for the Evaluation of Control Algorithms and Strategies*. Final Report for ASHRAE RP-825.

Haves, P., L.K. Norford, and M. DeSimone. 1998. A Standard Simulation Testbed for the Evaluation of Control Algorithms and Strategies. *ASHRAE Transactions* 104(1).

Hill, R.O. 1995. Applied Change of Mean Detection Techniques for HVAC Fault Detection and Diagnosis and Power Monitoring. M.S. thesis, Department of Architecture, Massachusetts Institute of Technology, Cambridge, Massachusetts.

House, J.M., H. Vaezi-Nejad, and J.M. Whitcomb. 2001. An Expert Rule Set for Fault Detection in Air-Handling Units. *ASHRAE Transactions* 107(1).

Hyvarinen, J., ed. 1996. *IEA Annex 25 Final Report*. VTT, Espoo, Finland.

Johansson, R. 1993. *System Modeling and Identification*. Prentice-Hall, Englewood Cliffs, New Jersey.

Krause, P.C. 1986. *Analysis of Electrical Machinery*. McGraw-Hill, New York.

Downloaded by [] at 07:55 19 July 2016

Lee, W.Y., C. Park, and G.E. Kelly. 1996a. Fault Detection of an Air-Handling Unit Using Residual and Parameter Identification Methods. *ASHRAE Transactions* 102(1):528-539.

Lee, W.Y., J.M. House, C. Park, and G.E. Kelly. 1996b. Fault Diagnosis of an Air-Handling Unit Using Artificial Neural Networks. *ASHRAE Transactions* 102(1):540-549.

Leeb, S.B. 1993. A Conjoint Pattern Recognition Approach to Nonintrusive Load Monitoring. Ph.D. dissertation, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, Massachusetts.

Leeb, S.B., S.R. Shaw, and J.L. Kirtley, Jr. 1995. Transient Event Detection in Spectral Envelope Estimates for Nonintrusive Load Monitoring. *IEEE Transactions on Power Delivery*. 10(3):1200-1210, July (also Paper 95 WM 042-2-PWRD IEEE PES Winter Meeting, New York, NY, January, 1995).

Li, X., H. Vaezi-Nejad, and J.-C. Visier. 1996. Development of a Fault Diagnosis Method for Heating Systems Using Neural Networks. *ASHRAE Transactions* 102(1):607-614.

Lorenzetti, D. and L.K. Norford. 1992. Measured Energy Consumption of Variable-Air-Volume Fans Under Inlet Vane and Variable-Speed Drive Control. *ASHRAE Transactions*. 98(2):371-379.

Luo, D., L.K. Norford, S.R. Shaw, and S.B. Leeb. 2002. Monitoring HVAC Equipment Electrical Loads from a Centralized Location—Methods and Field Test Results. Accepted for publication in *ASHRAE Transactions* and to be published in 108(1).

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 R.D. Little. 1993. Fault Detection and Load Monitoring in Ventilation Systems. *ASHRAE Transactions* 99(1):590-602.

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, and D. Luo. 2000. Demonstration of Fault Detection and Diagnosis Methods in a Real Building (ASHRAE 1020-RP). ASHRAE 1020-RP *Final Report*. ASHRAE, Atlanta. Report is posted on the TC4.11 website under www.ashrae.org.

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 for Air-Handling Units (ASHRAE 1020-RP). *International Journal of HVAC&R Research* 8(1).

Peitsman, H.C. and V.E. Bakker. 1996. Application of Black-Box Models to HVAC Systems for Fault Detection. *ASHRAE Transactions* 102(1):628-640.

Ruch, D. and D.E. Claridge. 1992. A Four-Parameter Change-Point Model for Predicting Energy Consumption in Commercial Buildings. *ASME Journal of Solar Energy Engineering* 114 (May):77-83.

Ruch, D., L. Chen, J.S. Haberl, and D.E. Claridge. 1993. A Change-Point Principal Component Analysis (CP/PCA) Method for Predicting Energy Usage in Commercial Buildings: The PCA Model. *ASME Journal of Solar Energy Engineering* 115 (May):77-84.

Salsbury, T.I. 1996. Fault Detection and Diagnosis in HVAC Systems Using Analytical Models. Ph.D. thesis, Loughborough University, U.K.

Seem, J.E., J.M. House, and C.J. Klaassen. 1998. Volume Matching Control: Leave the Outdoor Air Damper Wide Open. *ASHRAE Journal* 40(2):58-60.

Stylianou, M. and D. Nikanpour. 1996. Performance Monitoring, Fault Detection, and Diagnosis of Reciprocating Chillers. *ASHRAE Transactions* 102(1):615-627.

Tong, D. 1989. Intelligent and Healthy Buildings. *Proceedings of Designing for Environmental Quality 89*, Birmingham Polytechnic/RIBA, Solihull, U.K.

Tsutsui, H. and K. Kamimura. 1996. Chiller Condition Monitoring Using Topological Case-Based Modeling. *ASHRAE Transactions* 102(1):641-648.

Yoshida, H., T. Iwami, H. Yuzawa, and M. Susuki. 1996. Typical Faults of Air-Conditioning Systems and Fault Detection by ARX Model and Extended Kalman Filter. *ASHRAE Transactions* 102(1):557-564.

Downloaded by [] at 07:55 19 July 2016