
Automated Proactive Fault Isolation: A Key to Automated Commissioning

Srinivas Katipamula, PhD
Member ASHRAE

Michael Brambley, PhD
Member ASHRAE

ABSTRACT

Inadequate initial commissioning, the inability of the building operators to adequately monitor all building equipment and systems, and lack of proper maintenance or retro-commissioning lead to inefficient operations and reduced lifetimes of the heating, ventilating, air-conditioning (HVAC), and other energy-using equipment in buildings. Regularly scheduled manual maintenance and re-commissioning can help address these problems but can be labor intensive and perceived as expensive. Automated proactive commissioning and diagnostic technologies could help address two of the main barriers to commissioning and good maintenance—labor requirements and the costs associated with them. Automated proactive commissioning and diagnostic tools could reduce both the cost and time associated with commissioning, as well as enhance the persistence of commissioning fixes. Automation even offers the potential to go well beyond just monitoring and detecting faults to automatically correcting problems by compensating for sensor problems, reconfiguring controls, or changing control algorithms dynamically.

In this paper, we present a generic model for automated continuous commissioning and then delve in detail into one of the processes, proactive testing for fault isolation, which is key to automating commissioning. The automated commissioning process uses passive observation-based fault detection and diagnostic techniques, followed by automated proactive testing for fault isolation, automated fault evaluation, and automated reconfiguration of controls to continuously keep equipment controlled and running as intended. Only when hard failures occur or a physical replacement is required does the process require human intervention, and then sufficient

information is provided by the automated commissioning system to target manual maintenance where it is needed.

We then focus on fault isolation by presenting detailed logic that can be used to automatically isolate faults in valves, a common component in HVAC systems, as an example of how automated proactive fault isolation can be accomplished. Although automated passive fault detection and diagnostics have been tested in the field on actual equipment, the proactive process described in this paper has only been tested against simulations and has not yet been field tested. We conclude the paper with a discussion of how this approach to isolating faults can be applied to other common HVAC components and their automated commissioning, followed by a summary of key conclusions of the paper.

INTRODUCTION

Any commissioning provider, researcher, energy auditor, or mechanical service technician who has taken a moderately careful look at the heating, ventilating, and air-conditioning (HVAC) systems and equipment in any commercial building, new or existing, can tell stories about the dismal state of operation and maintenance (O&M). It is common to find economizers that don't modulate dampers, valves that leak, simultaneous heating and cooling because both heating and cooling valves are open, excessive use of reheat in terminal boxes during cooling because the setpoint for air leaving the air handlers is too low, and air-conditioning systems that are improperly charged and operate with dirty heat exchangers and filters (Katipamula et al. 2003a; Lunneberg 1999; Houghton 1997). It is also not uncommon for lights and other building systems to run 24 hours per day even though the building is unoccupied for several hours each day. These are a few of the

Srinivas Katipamula and Michael Brambley are staff scientists at Pacific Northwest National Laboratory, Richland, WA.

common conditions found that cause substantial energy waste in the commercial building stock. Although there are no reliable estimates of nationwide energy impacts associated with inefficient operations, there is a general consensus that the costs range from 10% to 30% (Ardehali and Smith 2002; Ardehali et al. 2003; Liu et al. 2002; Claridge et al. 1996).

Much of the energy waste in commercial buildings can be reduced by properly commissioning buildings and their systems. Building commissioning has been described and promoted in a number of publications (Haas and Sharp 1999; PEI 1997, 2000; DOE and GSA 1998; DOE 2002; Claridge et al. 2000; Liu et al. 2002, 2003). When executed effectively, commissioning and retro-commissioning eliminate such problems. This leads to energy savings, monetary savings on energy and peak electric demand, extended equipment life, and greater occupant comfort and satisfaction with indoor conditions. Energy savings reported for commissioning of existing buildings range from a few percent to over 60% with most reported savings in the range of 10% to 30% (Haas and Sharp 1999; Claridge et al. 2000; Liu et al. 2002).

Despite all the benefits, commissioning comes at a cost. A recent report (Quantum Consulting 2003, pp. 1–3) on commissioning in public buildings states, “While the concept of commissioning is increasingly accepted, there are still barriers—particularly with regard to cost—to implementation of the kind of thorough, independent third-party commissioning that is necessary for the full benefits of commissioning to be realized.” Only a small fraction of new construction and a very small fraction of existing buildings have been commissioned. Even when performed, pressures exist to keep costs down, which in some cases limits the depth to which the commissioning is performed. The authors hypothesize that costs play an important role in limiting the practice of commissioning for the building stock. Replacements or supplements to commissioning that reduce cost could, in the long run, better promote the objectives of commissioning. Key to this is reducing the labor intensity of commissioning by automating as many of the processes involved as possible. Compared to the cost of labor, automation technology is inexpensive.

To address this issue, a project was launched around 2000 with the overall goal of developing methods for improving the commissioning of HVAC systems through automation (PEI and Battelle 2003). Around the same time, another project was launched to investigate methods for automated diagnostics of selected building systems, including chillers, boilers, and chilled-water distribution systems (Sisk et al. 2003). Much of the work reported in this paper is derived from work performed by the authors on those two projects.

In this paper, we focus on describing methods for automatic proactive testing of common components of heating, ventilating, and air-conditioning (HVAC) systems for the purpose of fault isolation and diagnosis. We start by describing a generic process for automating commissioning to establish the context in which proactive testing would be used to support automated commissioning. This is followed by specification

of the logic used in proactive tests for isolating faults in valves, a common HVAC component, and a discussion of how this would fit into the process for automatically commissioning these and other system components.

GENERIC AUTOMATED COMMISSIONING PROCESS (AUTOCX)

Proactive automatic testing is used for fault isolation, one of the key functions in a process for automatically commissioning HVAC systems and their components. Figure 1 shows a generic process for automatically commissioning a building component or system. This process consists of five distinct functional processes: fault detection, fault isolation, fault evaluation, decision making, and corrective action.

In an automated system, fault detection would be done using one of the common methods (e.g., model based, rule based, case based, etc.) for automating fault detection and diagnostics (FDD). These methods are usually passive, automatically analyzing data on the current condition of the system as it operates and comparing that condition to the expected condition of the system, thereby detecting abnormal conditions or faults when the system is not operating correctly. Being passive, most applications of FDD observe conditions during ordinary operation; they do not initiate tests automatically to cause operation excursions. As a result, diagnostic systems must wait weeks, months, or even longer for changes in season before they experience a full range of operating conditions. As a result, their specificity in isolating faults to particular components and the depth of diagnosis are often limited, as is their ability to detect at any given time all faults present in a system. An example, would be the ability of an automated FDD tool for air handlers to detect that a fault exists in one of three temperature sensors, for the outdoor air, return air, or mixed air, but not to isolate which one of these sensors failed. Failure of one of these sensors can be detected easily from the thermodynamics of mixing airstreams (analytical redundancy), but isolation requires the

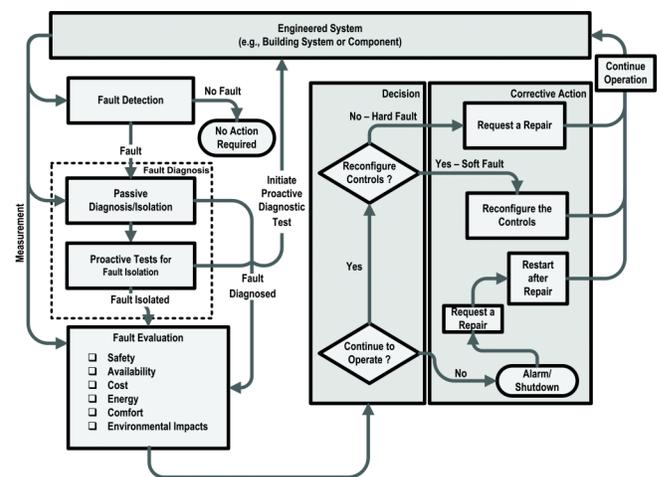


Figure 1 Diagram of a generic automated commissioning process.

ability to change (control) one additional physical parameter (Katipamula et al. 2003b).

Another process, proactive testing, can be used to introduce a controlled variable, providing the ability to further isolate the fault. When automated, this process provides the degree of fault isolation necessary to support automated commissioning. Proactive testing for fault isolation is the focus of this paper.

As shown in Figure 1, once the fault is adequately isolated, its impacts can be evaluated (the third process), providing information on changes in safety, system availability, energy use, costs, comfort, environmental impact, and other factors. This information then serves as input to deciding whether to take action and, if so, what action. Because of its importance to ensuring safety and minimizing damage to equipment, the first decision is whether to continue (or immediately stop) operation. If significant safety hazards or potential for equipment damage exists (based on the evaluation results), the system is stopped and an alarm is provided to the operator (owner or service provider) regarding the system fault, its seriousness, and suggested corrective actions needed before the system is restarted. Evaluation of impacts, deciding whether or not to shut down, and developing a corrective action plan could all be done manually, but by automating these processes, they can be done very quickly and continuously, 24 hours per day, seven days per week, year after year. Furthermore, the response time for shutting down equipment to minimize damage in most cases would be much shorter for an automated system. Ordinarily, when a system requires shut-down to prevent further damage, some sort of physical repair is necessary and, as a result, human intervention is required before the system can be restarted.

If a fault is detected, isolated, and evaluated but found to not present sufficient risk to shut down operation immediately, the next issue to address is whether the fault can be corrected simply by changing the control software code or values of parameters such as setpoints. These faults are sometimes called “soft” faults because they do not require physical repair of the system or cause the system to stop operating. Examples of soft faults include incorrect setpoints, biased sensors, incorrectly calibrated sensors, dampers, and valves, incorrect control code (e.g., logic errors), incorrectly tuned control algorithms, incorrect values for control-code parameters, and incorrect schedules. Because they are amenable to correction by software changes or changes in the values of parameters, these faults can be automatically corrected (see PECE and Battelle [2003] for an example of an automatic correction process). When corrected automatically, the commissioning loop is closed (see Figure 1) and the system continues to operate properly without interruption. We call this process *automated commissioning*. It performs all aspects of retro-commissioning a piece of equipment or system automatically, from determining whether it is configured correctly to correcting any faults found, then retesting by continuously monitoring and detecting any remaining or new faults in the system.

This process provides the foundation for continuous commissioning on the time scale of minutes and seconds rather than days, weeks, or months as when done manually.

When a fault has been found that cannot be corrected by reconfiguring the controls, a decision must be made whether to tolerate the fault or correct it. In cases where a fault is judged to have sufficiently low impact to be tolerable, operation may continue without correction of the problem by repair. This might be the case, for example, when an electric motor is found to have dropped by 5% in efficiency. Such a change might be judged sufficiently small to continue operation unless and until the efficiency has decreased more (e.g., by 10% or 15%). This decision could be based on algorithms that examine the cost-effectiveness of taking the motor out of operation and repairing it versus losing some money on less efficient operation in the present condition. Similarly, a threshold could be established based on costs to take the motor out of operation and repair it when the threshold is exceeded. With an automated system, this decision is constantly revisited in real time and a decision to repair made when the degradation can no longer be tolerated. In that case, the equipment is repaired and then placed back into operation. By performing all key processes automatically, except the repair, systems would be kept in optimal operating condition at all times, rather than going unnoticed potentially for long periods of time as is common in buildings today.

We call the process described above, and shown schematically in Figure 1, *automated proactive commissioning*, even though parts of it are not completely automated. When “soft” faults are found and automatically corrected, the process is fully automated. When a “hard” fault occurs and physical repair is required, the process becomes semi-automated commissioning. In this latter case, all parts of the commissioning except for repair are automated.

The proactive diagnostic process can help in diagnosing and isolating faulty operations to a much greater extent than passive diagnostics, but it is intrusive in nature. Some building owners and operators may consider this to be disruptive to the normal operation of their building systems. They may not, however, if such proactive tests can be conducted quickly enough so that acceptable control of the building systems is maintained. Alternatively, entirely proactive commissioning procedures could provide “continuous” commissioning if they were periodically triggered (e.g., once a day, week, or perhaps a month). These procedures might be scheduled to occur during unoccupied hours to further reduce their intrusion on normal operations. Because proactive tests are potentially disruptive, the criteria and thresholds to activate them should be thoroughly analyzed before implementation. If the criteria and thresholds are loosely defined and the proactive tests are initiated too often during occupied times, they could potentially be deemed unacceptably disruptive by building owners, operators, or occupants.

The remainder of this paper focuses on the logic required to automate one of the key parts of the overall automated

commissioning process, proactive testing for fault isolation. The logic is developed for one of the common devices used in HVAC systems, valves. For more information on automated correction of “soft” faults and other applications, see PECEI and Battelle (2003) and Katipamula et al. (2003b).

BACKGROUND INFORMATION ON VALVES

In HVAC applications, valves are widely used to control fluid flow to heat exchangers (the coils) in air-handling units (AHUs) and to primary and secondary distribution systems. The valves can be two-way or three-way depending on requirements. In an AHU, the zone thermostat activates the valves (i.e., selects cooling or heating mode), but feedback of supply-air temperature measurements is used to control the amount of valve opening (by sending a control signal to the valve actuator). Two common valve fault modes are sticking and leaking. Stuck valves manifest their effects through lack of control of supply-air temperature, which directly impacts the comfort of occupants, while leaky valves are difficult to detect unless the leak is severe. A cooling-coil valve leak equal to 20% of full flow may only produce a 1°C (1.8°F) drop in air temperature across the coil (Bushby et al. 2002). This is roughly the accuracy of many commercial air-temperature sensors. As a result, smaller leaks could be masked by the inability to reliably measure an impact.

In most AHUs, the temperature of the air between the heating and the cooling coils is not measured. In order to detect heating valve leakage under all conditions in the cooling mode, this air temperature needs to be measured and an additional air-temperature sensor between the two coils is required. In the absence of this measurement, the method can still detect many instances of leaking valves. Some cases, however, will be missed using this method when the heat gain and heat loss from leaks by both the heating and cooling valves exactly balance each other.

GENERAL APPROACH FOR IDENTIFYING FAULTY AND MALFUNCTIONING VALVES

In this section we present a methodology to isolate stuck and leaky valves with passive and proactive diagnostic tests. Several researchers have developed passive FDD methods for detecting faulty operations including leaky cooling and heating valves (Haves et al. 1996; Lee et al. 1996a, 1996b, 1997; Peitsman and Soethout 1997; Ngo and Dexter 1999; Dexter and Ngo 2001; Norford et al. 2002).

One of the main symptoms of a leaky or stuck valve is the inability of the AHU to maintain the supply air temperature. Most methods for detecting these faults are based on passive monitoring of operation and collection of data from which to determine the current state of a system. The current state is then compared to the expected state under identical driving conditions to determine whether a fault exists or not. A deviation judged sufficient results in an alarm, which indicates that a fault has been detected. The alarm may also include information on the type of fault, the nature of the fault, and often

specification to some degree of where in the system the fault has occurred. We call such methods *passive fault detection* or *passive fault detection and diagnostics*. Detailed descriptions of methods for passive FDD for air handling are provided in PECEI and Battelle (2003), Schein et al. (2003), and House et al. (2001). Furthermore, valve and damper FDD are described in Section 7.10 of PECEI and Battelle (2003).

The detailed logic for isolating faults in heating and cooling valves in HVAC systems presented in this paper starts after a valve fault has been detected using passive FDD but before the fault has been sufficiently isolated to determine its impacts well or select a corrective action. This is the second part of the second major function in the automated commissioning process shown in Figure 1.

Chilled-Water Valve Fault Detection and Diagnostics (FDD)

The process for isolating a chilled-water (CW) valve fault is presented in the flowcharts in Figure 2 through Figure 6. These flowcharts and all others in this paper use symbology, where a circle represents a connection to or from another process or subprocess, diamonds represent decisions, rectangular boxes represent actions, and rounded boxes represent end states or conclusions. A passive FDD process that identifies an AHU as unable to maintain the supply-air setpoint¹ precedes Figure 2 (as indicated by the circle above the process boundary). Logic for this passive process can be found in PECEI and Battelle (2003).

Isolation of the chilled-water valve fault is divided into five subprocesses:

- Initial CW valve diagnostics (Figure 2)
- Process to detect and isolate a CW valve stuck fully open (Figure 3)
- Process to detect and isolate a CW valve stuck partially open (Figure 4)
- Process to quantify a hot water (HW) valve leak while the AHU is in cooling mode (Figure 5)
- Process to detect and isolate a CW valve stuck fully closed (Figure 6)

The connections between these subprocesses are shown by circles on the flowcharts. The measured variables needed for the CW valve diagnostics are identified in Table 1.

The initial diagnostic process (Figure 2) detects the presence of any of four faults—simultaneous heating and cooling by the heating and cooling coils, a CW supply temperature that is too high, potential problem with the control logic for the chilled-water valve, or a leaking HW valve while the AHU is in cooling mode. If none of these conditions is found, the flowchart leads to a connection to one of the additional five subprocesses, the

¹ If supply-air temperature is measured downstream of the supply-air fan, the process adjusts the temperature to reflect the heat gain from the fan.

Table 1. Measured Variables Required by the Proactive Diagnostic Process for Chilled-Water Valves

Measured Variables
Supply-air temperature
Mixed-air temperature
Supply-air temperature setpoint
Air temperature between the heating and cooling coils
CW supply temperature
CW temperature setpoint
CW valve signal
Heating valve status (ON/OFF)
Supply fan signal (if variable speed)

specific process depending on conditions (i.e., the branch of the flowchart) found while executing the analytic tests shown in Figure 2. The reader should also note that the diagnostics shown in Figure 2 are only undertaken if preceding passive diagnostics found that while in cooling mode the AHU was unable to maintain the supply-air temperature at its setpoint.

All tests depicted in Figure 2 are passive (referred to as analytic tests) except those shown in the box labeled “Proactive Diagnostic Process 1—CW Control Logic Problem,” which requires a physical test in which automatic control of the CW valve is overridden and the CW valve signal is proactively set to 100% open. This test enables identification of potential problems with the CW valve control logic. The subprocesses in Figure 3 through Figure 6 all similarly include proactive physical tests.

The process shown in Figure 3 tests whether the CW valve is stuck fully open (at 100% flow). It identifies any of the following problems if they are present:

- CW valve control logic or controller hardware problem
- CW valve stuck fully open
- CW valve modulating, but not properly

All branches of the flowchart in Figure 3 lead to end states corresponding to isolation of a fault.

Figure 4 shows the subprocess that tests whether the CW valve is stuck partially open and leads to one of three possible faults:

- CW valve stuck partially open
- CW valve not modulating properly or the valve may be faulty
- CW valve is closing fully but not opening fully, the cooling coil may be fouled, or the valve hardware is faulty.

As with Figure 3, all branches of the chart in Figure 4 lead to end states.

The subprocess shown in Figure 5 is performed when a hot-water valve leak is found while executing the diagnostics

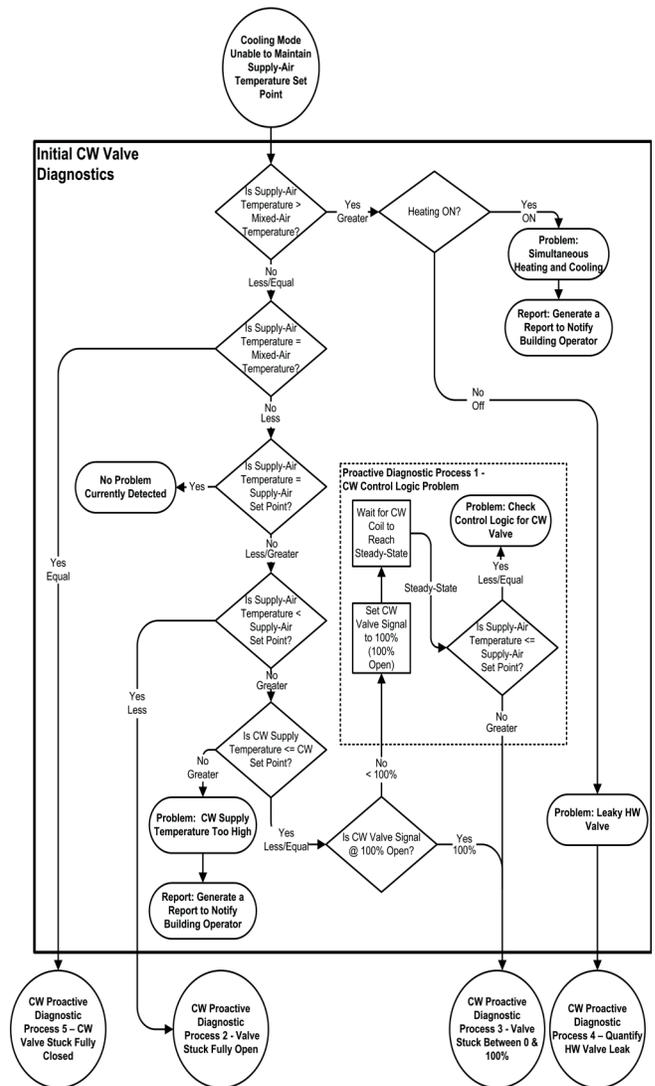


Figure 2 Flowchart for initial chilled-water valve diagnostics.

shown in Figure 2. It represents a test to quantify the hot-water valve leak. To estimate the amount of leakage, a lookup table is required. The lookup table can be developed using the on-line training method described in PECI and Battelle (2003, pp. 129–130). By comparing the measured temperature difference across the heating coil to the values in the table, an estimate of the leakage is obtained.

Figure 6 shows the subprocess to determine whether the chilled-water valve is stuck fully closed. It leads to identification of faulty CW valve control (logic or controller hardware) and identification of the CW valve stuck fully closed or directs the process to continue with the subprocess in Figure 4, which checks whether the CW valve is stuck at a partially open position as described previously. The process in Figure 4 leads only to end states, terminating the CW diagnostic process and identifying one of the end states as existing.

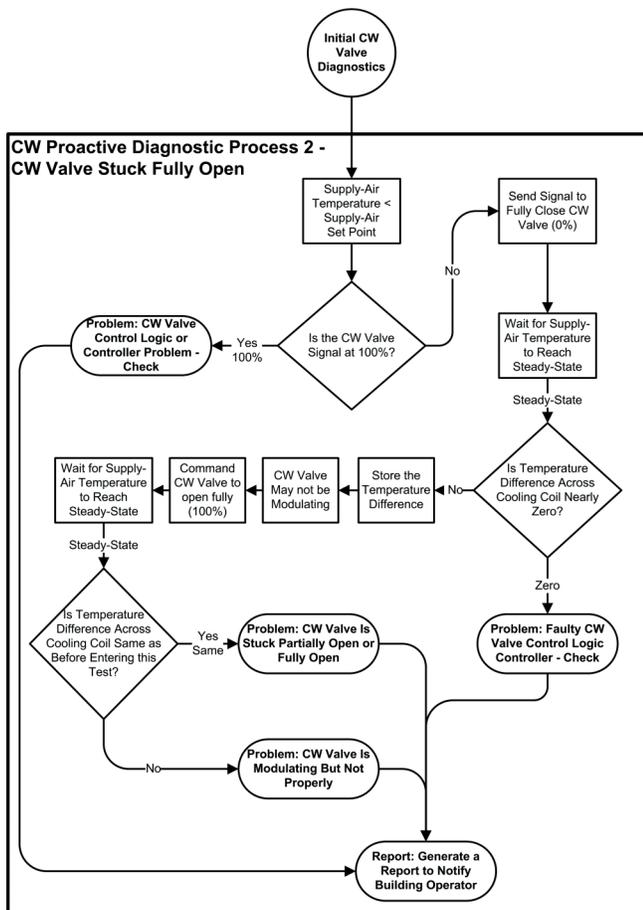


Figure 3 Proactive diagnostic process to detect a chilled-water valve stuck fully open.

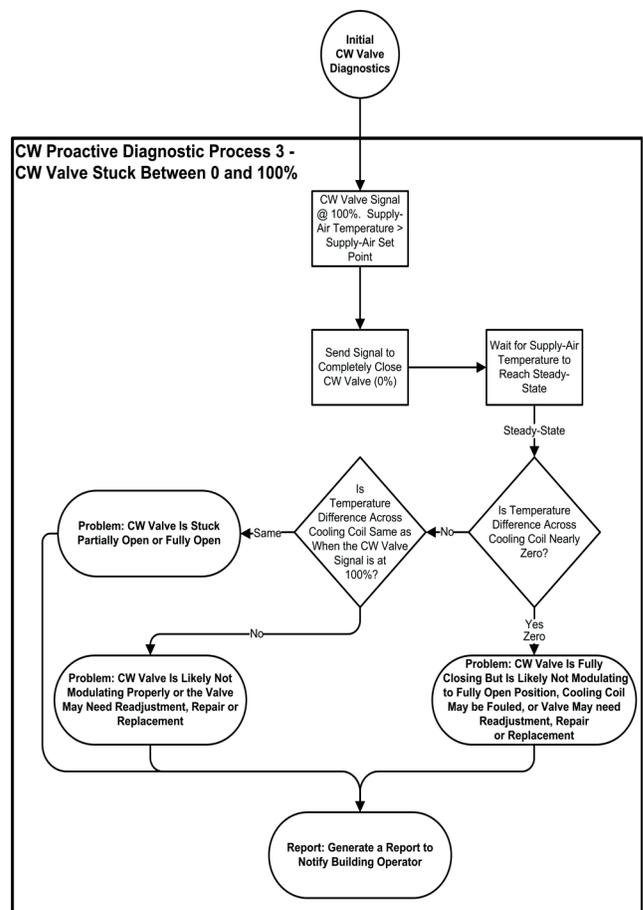


Figure 4 Proactive diagnostic process to detect a chilled-water valve stuck partially open.

Approach to Hot-Water Valve FDD

The measured variables needed to identify malfunctioning and faulty hot-water valves are identified in Table 2, and the process is shown in Figure 7 through Figure 11. The approach is similar to that for detecting faults with a chilled-water valve described in the previous section. The full process is divided into five subprocesses:

- Initial HW valve diagnostics (Figure 7)
- Process to detect and isolate an HW valve stuck fully open (Figure 8)
- Process to detect and isolate an HW valve stuck partially open (Figure 9)
- Process to quantify a chilled-water (CW) valve leak while the AHU is in heating mode (Figure 10)
- Process to detect and isolate an HW valve stuck fully closed (Figure 11)

The initial diagnostic process (Figure 7) detects the presence of any of four faults: simultaneous heating and cooling by the heating and cooling coils, a HW supply temperature that is

too low, a potential problem with the control logic for the hot water valve, or a leaking CW valve while the AHU is in heating mode. If none of these conditions is found, the flowchart leads to a connection to one of the other five subprocesses, the specific process depending on conditions found while executing the analytic tests shown in Figure 7. The reader should also note that the diagnostics shown in Figure 7 are only undertaken if preceding passive diagnostics find that while in heating mode the AHU is unable to maintain the supply-air temperature at its setpoint.

The process shown in Figure 8 tests whether the HW valve is stuck fully open. It identifies any of the following problems if they are present:

- HW valve control logic or controller hardware problem
- HW valve stuck fully open
- HW valve modulating, but not properly

All branches of the flowchart in Figure 8 lead to end states corresponding to isolation of a fault, so if the diagnostic process leads to the flowchart in Figure 8, an HW valve fault will be identified.

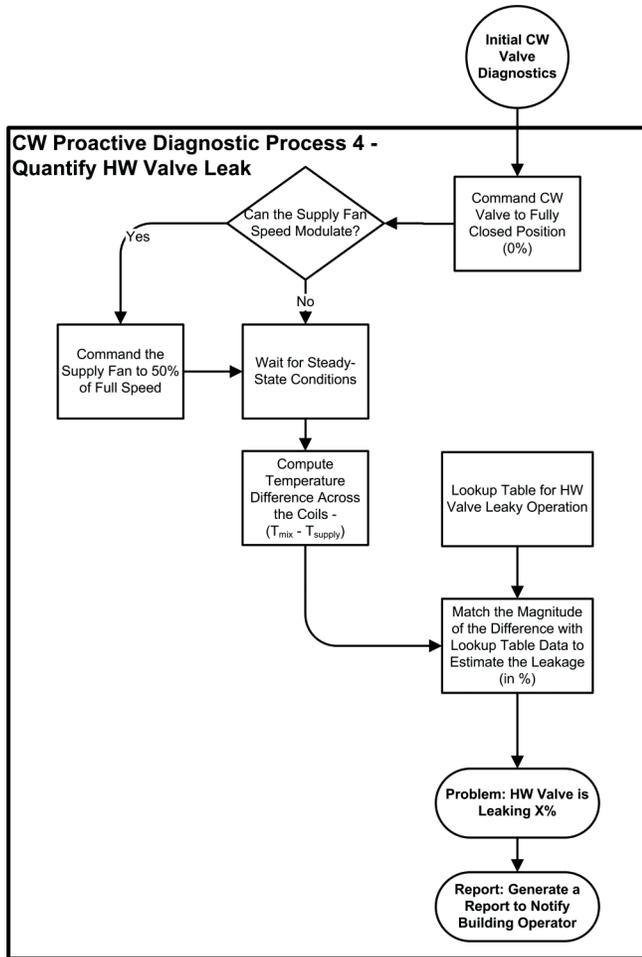


Figure 5 Proactive diagnostic process to quantify a hot-water valve leak while in cooling mode.

Figure 9 shows the subprocess that tests whether the HW valve is stuck partially open and leads to one of three possible faults:

- HW valve stuck partially open
- HW valve not modulating properly, or the valve itself may be faulty
- HW valve is closing fully but not opening fully, the heating coil may be fouled, or the HW valve itself is faulty

As with the flowchart in Figure 8, all branches of the chart in Figure 9 lead to end states.

The subprocess shown in Figure 10 depicts tests to determine whether the CW valve is leaking while the AHU is in heating mode. The subprocess shown in Figure 10 is performed when a leaky chilled-water valve is found while executing the diagnostics shown in Figure 7. It represents a test to quantify the chilled-water valve leak. To estimate the amount of leakage, a lookup table is required. As with the parallel test for HW valve leakage in Figure 5, this lookup

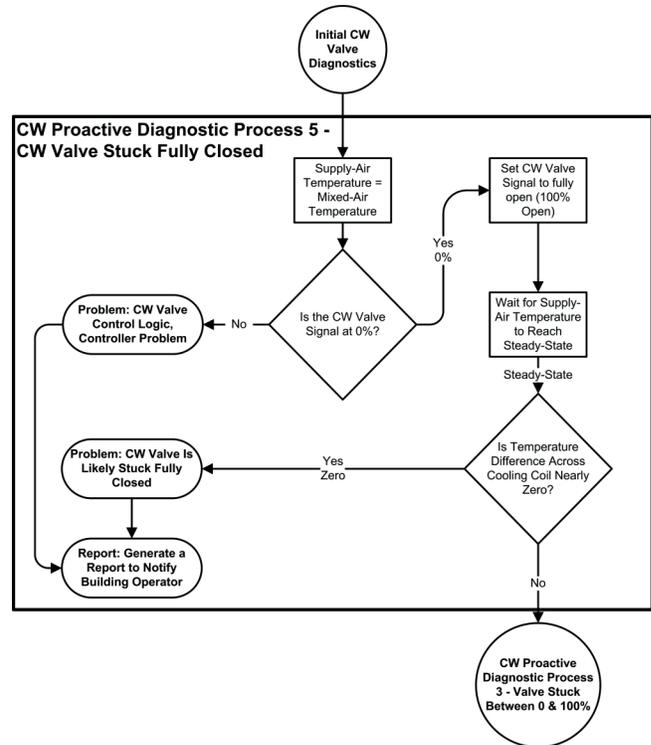


Figure 6 Proactive diagnostic process to detect a chilled-water valve stuck fully closed.

Table 2. Measured Variables Required by the Proactive Diagnostic Process for Hot Water Valves

Measured Variables
Supply-air temperature
Mixed-air temperature
Supply-air temperature setpoint
Air temperature between the heating and cooling coils
HW supply temperature
HW temperature setpoint
HW valve signal
Cooling valve status (ON/OFF)
Supply fan signal (if variable speed)

table can be developed using the on-line training method described in PECE and Battelle (2003). By comparing the measured temperature difference across the cooling coil to the values in the table, an estimate of the leakage is obtained.

Figure 11 shows the subprocess to determine whether the HW valve is stuck fully closed. It leads to identification of faulty HW valve control (logic or controller hardware) and identification of the HW valve stuck fully closed or directs the

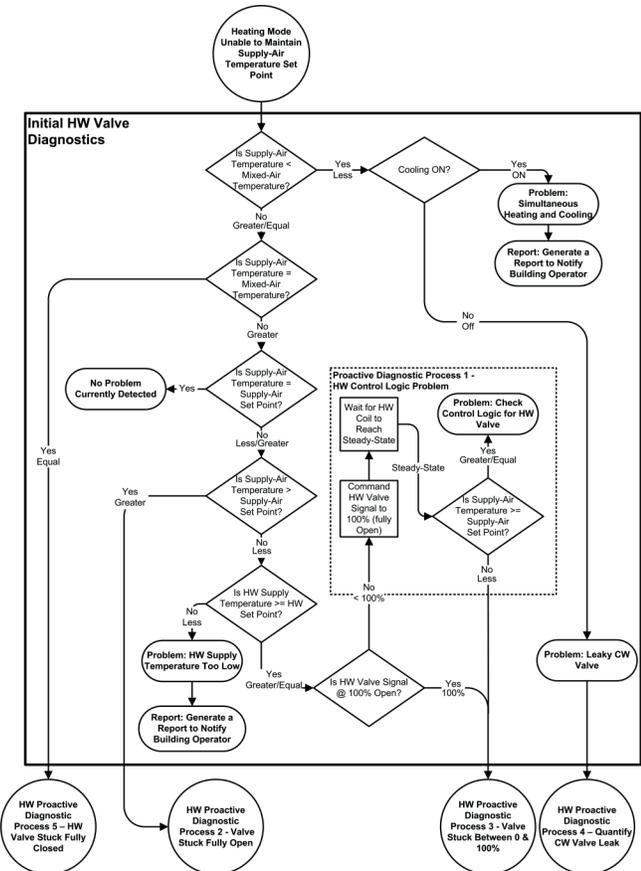


Figure 7 Flowchart for initial hot-water-valve diagnostics.

process to continue with the subprocess in Figure 9, which checks whether the HW valve is stuck at a partially open position as described previously. The process in Figure 9 leads only to end states, terminating the HW diagnostic process and identifying one of the end states as corresponding to the current state of the HW valve.

DISCUSSION

The logic provided in the preceding sections provides a critical component for automating the commissioning of valves in HVAC systems. Although valves are but one component in HVAC systems, the process illustrated can, by reasoning about physical behavior, be extended to other HVAC components (see PECE and Battelle [2003]). The logic presented in this paper focuses on adequately isolating faults so that they can be evaluated with respect to impacts on energy use, cost, safety, etc., and specific actions can be identified and implemented to correct them. As shown in Figure 1, proactive tests to isolate faults are preceded by passive observational fault detection and isolation techniques (described, for example, in PECE and Battelle [2003]). When automated together, these five processes—fault detection, fault isolation, fault evaluation, and decision and implementation of corrective

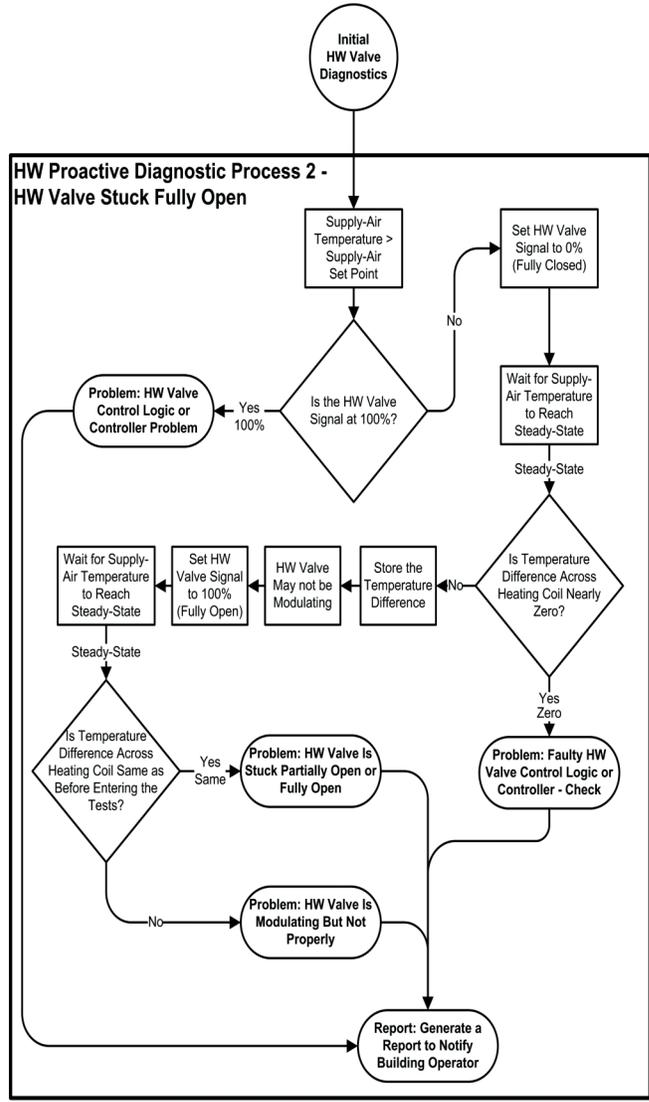


Figure 8 Proactive diagnostic process to detect a hot-water-valve stuck fully open.

action—form the foundation for automated commissioning of equipment and systems.

The logic presented in this paper for valves is easily automated in software as sets of if-then-else rules. Coupled with a continuous flow of data from a building automation system or separate data acquisition system, a system can be created for automatically and continuously commissioning HVAC components and systems. When a fault cannot be corrected or compensated for automatically, support can be provided by providing information to operation and maintenance personnel that enables them to quickly assess impacts and to plan and take corrective actions. In such cases, the process becomes semi-automated but much more streamlined than the one-time or periodic process generally used in commissioning and retro-commissioning today.

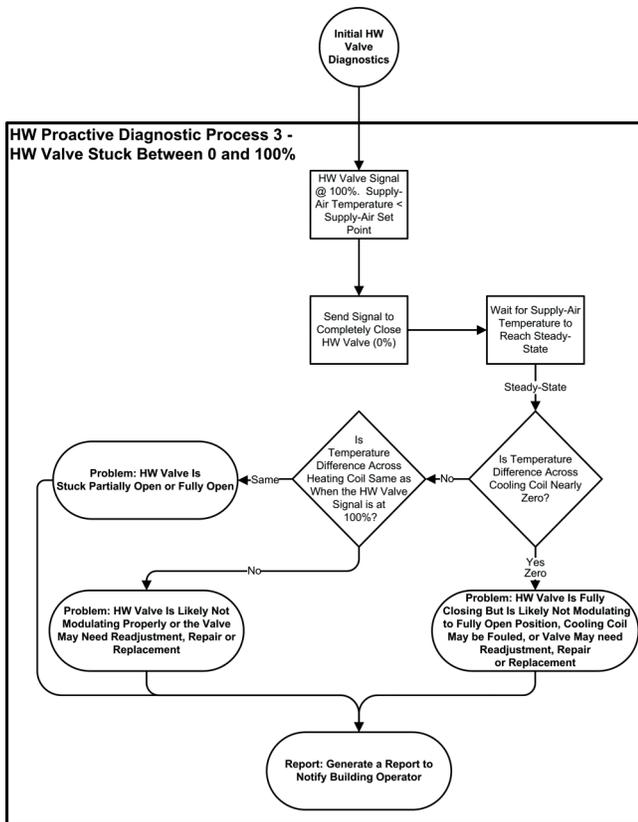


Figure 9 Proactive diagnostic process to detect a hot-water valve stuck partially open.

Although the logic is presented as deterministic in this paper, practical implementation involves adequately accounting for uncertainties in measured quantities when implementing it in software. If uncertainties are not adequately considered, this method will likely provide an unacceptably large number of false alarms. One method for including uncertainty involves establishing a tolerance band around each measured value. These tolerances are used to account for random noise, measurement uncertainty, and systematic bias in measurements (e.g., measurements that are consistently high or low relative to the true value of the measured variable). The tolerances are then propagated through all calculations and comparison tests. For example, to test if the outdoor-air temperature is greater than the return-air temperature at a specific time, not only should the outdoor-air temperature value be greater than the return-air temperature, it should be greater than the return-air temperature plus a multiple of the uncertainty of the difference between the two measured values in order to minimize the probability that the true outdoor air is actually less than or equal to the return-air temperature. The uncertainty of the difference between two measured variables is equal to the sum of the uncertainties for each of the two variables, e.g., to test whether the outdoor-air temperature (T_{out}) is equal to the return-air temperature (T_{ret}), if the tolerances on

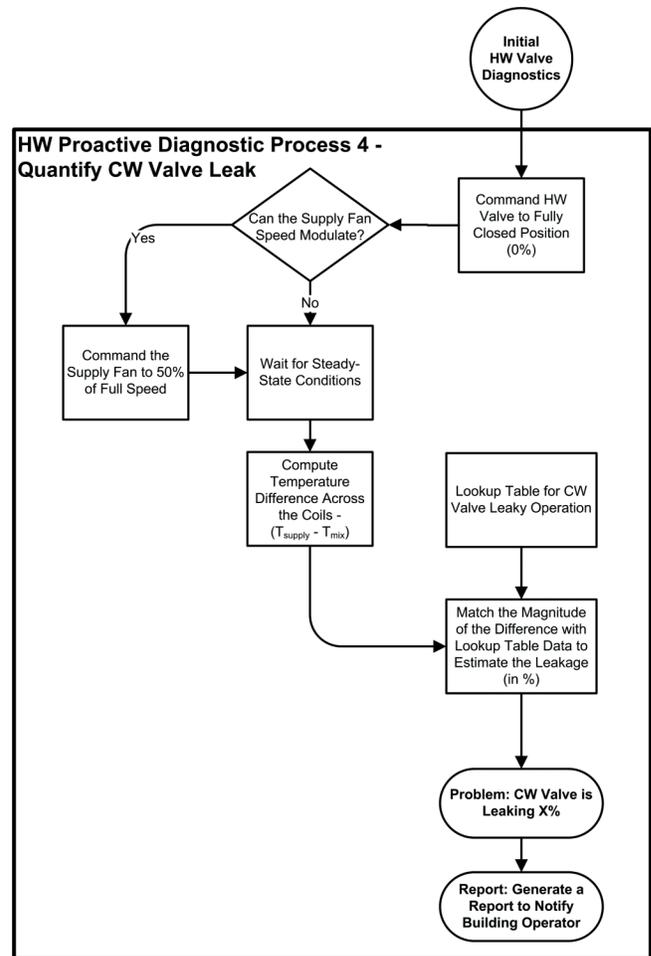


Figure 10 Proactive diagnostic process to detect and diagnose a leaky chilled-water valve while in heating mode.

both temperature measurements are set to $\pm 0.5^\circ\text{F}$, the outdoor-air temperature is declared equal to the return-air temperature only if the following condition is true:

$$|(T_{out} - T_{ret})| \leq 1.0$$

$$0.5 - (-0.5) = 1.0$$

The less-than, greater-than, less-than-or-equal-to, and greater-than-or-equal-to are constructed similarly using the assigned tolerances. Similarly, the uncertainty associated with other algebraic combinations of measured variables and tests can be evaluated using standard formulas for the propagation of errors in calculations (see, for example, NIST/SEMATECH 2005).

Small tolerances generally result in greater detection sensitivity and greater false alarm rates because smaller differences result in conditions being more easily satisfied. Actual implementation requires specifying tolerances for all variables and propagating them throughout the decision tree. Specifying such a large number of tolerances can present an unwieldy burden for users of fault detection and diagnosis

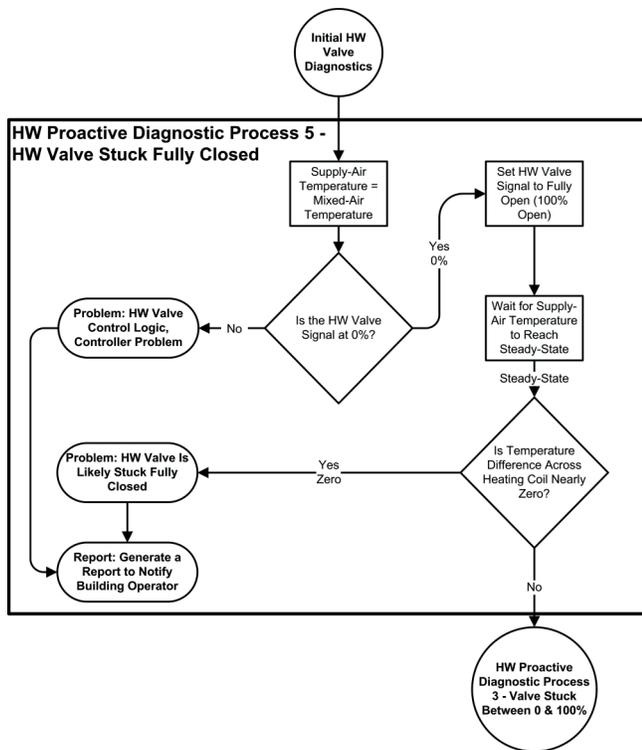


Figure 11 Proactive diagnostic process to detect a hot-water valve stuck fully closed.

(FDD) software; therefore, we recommend that developers of software using this method package sets of tolerance values for all variables such that each set represents a different level of detection sensitivity (and associated false alarm rate) for the entire diagnostics process. By adjusting tolerances while testing the performance of the resulting algorithms with actual measured data, developers can tune packages of tolerances and detection and diagnostic thresholds for all variables and rules to produce detection and diagnostic sensitivities varying from very low to very high, (e.g., very low, low, medium, high, and very high).² Users of the software then can start their use of the software using some intermediate sensitivity level (say, low) and then adjust the sensitivity up (e.g., to medium) or down (to very low) as experience is gained with performance of the FDD tool. This process is much like a radio listener adjusting the volume control on a radio (but with a small number of discreet sensitivity settings rather than the continuous variability of the volume control). There is generally no quantitative indication of volume level on the volume control of a radio, but by listening, adjusting the volume control, and obtaining feedback by listening at the new volume setting, the

² In principle, this approach to handling tolerances on variables could become unwieldy for the tool developer, as the number of rules, and, therefore, number of variables, increases significantly. The authors have not yet encountered a case in their work on HVAC diagnostics where that has occurred.

listener can adjust the control until the desired volume level is reached. The feedback process is essentially instantaneous for the radio listener and possibly days or even months for the user of FDD software, but the process is similar. No direct quantitative measure of detection sensitivity is necessary. The user simply builds experience with the detection sensitivity and adjusts the sensitivity upward or downward to achieve a desired resultant detection sensitivity and error rates.

CONCLUSIONS

A significant amount of information is generated, collected, and used during the commissioning process to ensure systems and buildings are designed, constructed, and operated in a manner that meets the owner's requirements. During retro-commissioning, functional tests are run to determine whether equipment and systems operate properly. Then, adjustments and repairs are made to bring the equipment into conformance with expectations. The development of automated continuous commissioning tools will reduce the time and cost associated with commissioning and retro-commissioning of HVAC systems and enhance system performance over the life of a building. The current research set out to identify and develop key enabling components essential for creation of automated commissioning tools. Five key functional processes were identified:

- Fault detection
- Fault isolation
- Fault evaluation
- Decision making
- Fault correction.

In automated commissioning, all five functional processes are automated. When any of these processes is automated and used in support of commissioning, the resulting overall process can be considered semi-automated commissioning. In the long run, after methods for all of the functional processes are fully developed, semi-automated commissioning likely will involve automation of the first three, followed by manual decision making and fault correction for hard faults with the response to soft faults being fully automated.

This paper focuses primarily on one of the key functional processes, fault isolation, to the level necessary to support automated correction of faults by proactive automatic testing. Logic is given specifically for the example of heating and cooling valves. Following specification of this logic, the handling of measurement uncertainty in software implementation of the logic by assignment and propagation of tolerances for measured variables is discussed along with its relationship to fault detection sensitivity and detection errors. We then propose a grouping of measurement tolerances into sensitivity categories (e.g., very low, low, normal, high, and very high—or low, medium, and high) as a practical solution for tolerance setting in software that uses these methods. Users can then set the sensitivity of the fault detection and isolation system simply by empirically adjusting the sensitivity much like a

music listener adjusts the volume on a stereo system (but with a longer characteristic time for feedback).

The paper also provides references for advancement on automation of other parts of the commissioning process.

ACKNOWLEDGMENTS

The work reported in this paper was funded partially by the American Air-Conditioning and Refrigeration and Technology Institute as part of the 21-CR program, the California Energy Commission as part of a California Public Interest Energy Research program managed by Architectural Energy Corporation, and the Building Technologies Program of the U.S. Department of Energy. Pacific Northwest National Laboratory is operated for the U.S. Department of Energy by Battelle Memorial Institute under contract DE-AC05-76RL01830. The authors also would like to thank Peter Armstrong for providing valuable comments on the manuscript and Susan Arey for editorial support.

REFERENCES

- Ardehali, M.M., and T.F. Smith. 2002. Literature review to identify existing case studies of controls-related energy-inefficiencies in buildings. Technical Report: ME-TFS-01-007, Department of Mechanical and Industrial Engineering, The University of Iowa, Iowa City, Iowa.
- Ardehali, M.M., T.F. Smith, J.M. House, and C.J. Klaassen. 2003. Building energy use and control problems: An assessment of case studies. *ASHRAE Transactions* 109(2):111–21.
- Bushby, S.T., N.S. Castro, J. Schein, and J. House. 2002. Using the virtual cybernetic building testbed and FDD test shell for FDD tool development. NISTIR 6818, National Institute of Standards and Technology, Gaithersburg, MD.
- Claridge, D.E., M. Liu, Y. Zhu, M. Abbas, A. Athar, and J.S. Haberl. 1996. Implementation of continuous commissioning in the Texas LoanSTAR program: Can you achieve 150% estimated retrofit savings revisited. *Proceedings of the 1996 ACEEE Summer Study on Energy Efficiency in Buildings*, pp. 4.59–4.67.
- Claridge, D.E., C.H. Culp, M. Liu, S. Deng, W.D. Turner, and J.S. Haberl. 2000. Campus-Wide Continuous CommissioningSM of University Buildings. *Proceedings of the 2000 ACEEE Summer Study on Energy Efficiency in Buildings*, pp. 3.101–3.112.
- Dexter, A.L., and D. Ngo. 2001. Fault diagnosis in air-conditioning systems: A multi-step fuzzy model-based approach. *HVAC&R Research* 7(1):83–102.
- Haasl, T., and T. Sharp. 1999. *A Practical Guide for Commissioning Existing Building*. Portland, OR: Portland Energy Conservation, Inc., and Oak Ridge, TN: Oak Ridge National Laboratory.
- Haves, P., T. Salisbury, and J.A. Wright. 1996. Condition monitoring in HVAC subsystems using first principles. *ASHRAE Transactions* 102(1):519–27.
- Houghton, D. 1997. Operating and maintaining rooftop air conditioners. *ASHRAE Journal* 39(12):50.
- 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):858–71.
- Katipamula, S., M.R. Brambley, N.N. Bauman, and R.G. Pratt. 2003a. Enhancing building operations through automated diagnostics: Field test results. *Proceedings of the 2003 Third International Conference For Enhanced Building Operations, October 13–15, 2003, Berkeley, California* [CD].
- Katipamula, S., M.R. Brambley, and L. Luskay. 2003b. Automated proactive techniques for commissioning air-handling units. *ASME Journal of Solar Energy Engineering* (special issue on emerging trends in building design, diagnosis and operation), *Transactions of the ASME* 125(1):282–91.
- Lee, W.Y., C. Park, and G.E. Kelly. 1996a. Fault detection of an air-handling unit using residual and recursive parameter identification methods. *ASHRAE Transactions* 102(1):528–39.
- 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):540–49.
- Lee, W.Y., J.M. House, and D.R. Shin. 1997. Fault detection of an air-handling unit using residual and recursive parameter identification methods. *ASHRAE Transactions* 102(1):528–39.
- Liu, M., D.E. Claridge, and W.D. Turner. 2002. *Continuous CommissioningSM Guidebook for Federal Energy Managers*. Washington, DC: Federal Energy Management Program, U.S. Department of Energy. www.eere.energy.gov/femp/pdfs/ccg01_covers.pdf.
- Liu, M., D.E. Claridge, and W.D. Turner. 2003. Continuous CommissioningSM of building energy systems. *ASME Journal of Solar Energy Engineering* (special issue on emerging trends in building design, diagnosis and operation), *Transactions of the ASME* 125(3):275–81.
- Lunneberg, T. 1999. When good economizers go bad. E Source Report ER-99-14, E Source, Boulder, CO.
- Ngo, D., and A.L. Dexter. 1999. A robust model-based approach to diagnosing faults in air-handling units. *ASHRAE Transactions* 105(1):1078–86.
- NIST/SEMATECH. 2005. *NIST/SEMATECH e-Handbook of Statistical Methods*, Section 2.5.5. Gaithersburg, MD: National Institute of Standards and Technology. www.itl.nist.gov/div898/handbook/.
- Norford, L.K., J.A. Wright, R.A. Buswell, D. Luo, C.J. Klaassen, and A. Suby. 2002. Demonstration of fault detection and diagnosis methods for air-handling units (ASHRAE RP-1020). *HVAC&R Research* 8(1):41–71.
- Peitsman, H.C., and L.L. Soethout. 1997. ARX models and real-time model-based diagnosis. *ASHRAE Transactions* 103(1):657–71.

- PECI. 1997. Commissioning for better buildings in Oregon. Salem, OR: Oregon Energy Office. Portland Energy Conservation Inc.
- PECI. 2000. *New Construction Commissioning Handbook for Facility Managers*. Salem, OR: Oregon Energy Office. Portland Energy Conservation Inc.
- PECI and Battelle. 2003. Methods for automated and continuous commissioning of building systems. Final report. ARTI-21CR/610-30040-01, Air-Conditioning & Refrigeration Technology Institute, Washington, DC. Portland Energy Conservation, Inc. and Battelle Northwest Division. Available at www.arti-21cr.org/research/completed/finalreports/30040-final.pdf.
- Quantum Consulting. 2003. *Market Progress Evaluation Report—Commissioning in Public Buildings Project, No. 3*, Report #E03-107. Portland, OR: Northwest Energy Efficiency Alliance.
- Schein, J., S.T. Bushby, N.S. Castro, and J.M. House. 2003. Results from field testing of air handling unit and variable air volume box fault detection tools, NISTIR 6994, National Institute of Standards and Technology, Gaithersburg, MD.
- Sisk, D.R., M.R. Brambley, T.A. Carlon, and R.S. Briggs. 2003. *Automated Diagnostics Software Requirements Specification, Version 1.1*. Battelle Report PNWD-3316. Battelle Northwest Division, Richland, WA. Available at www.archenergy.com/cec-eeb/P2-Diagnostics/P2-Diagnostics_Reports/P5/D-2.5.4a-SoftwareRequirementsSpecification1-1pnwd3316.pdf.
- DOE and GSA. 1998. *Building Commissioning Guide, Version 2.2*. Washington, DC: U.S. Department of Energy, Federal Energy Management Program (FEMP) and U.S. General Services Administration, Public Buildings Service.
- DOE. 2002. *National Best Practices Manual Commissioning*. Washington, DC: U.S. Department of Energy.

DISCUSSION

Xiaohui Zhou, Assistant Scientist, Iowa Energy Center, Ankeny, IA: In the automated commissioning process, is there a difference in which subsystem function test is done first?

Srinivas Katipamula: The automated commissioning process follows a predefined process (based on first principles) until an end state is reached. So, the automated functional tests are performed as required by the process. In other words, the functional tests are not independent.