Automated Fault Diagnostics for AHU-VAV Systems: A Bayesian Network Approach

Adam Regnier
Drexel University, United States of America, ar626@drexel.edu

Jin Wen
Drexel University, United States of America, jinwen@drexel.edu

Automated Fault Diagnostics for AHU-VAV Systems: A Bayesian Network Approach

Adam REGNIER[1]*, Jin WEN[2]

[1]Drexel University,
Department of Civil, Architectural, and Environmental Engineering,
Philadelphia, PA, USA
ar626@drexel.edu

[2] Drexel University,
Department of Civil, Architectural, and Environmental Engineering,
Philadelphia, PA, USA
jinwen@drexel.edu

* Corresponding Author

ABSTRACT

Although it is widely accepted that 20%-30% of total HVAC energy in commercial buildings is wasted due to faulty or inefficient operation, there presently exist no widely adopted solutions to identify and remedy this waste. This lack of adoption is due primarily to the high upfront costs associated with the manual process of customizing commissioning solutions for each individual building. In order to reduce these costs, diagnostic technologies must automate the process of installing and customizing solutions for each implementation. A novel approach to addressing this problem for air handling units (AHUs) and variable air volume (VAV) boxes is presented here. This strategy utilizes a Bayesian network to identify and understand the system operation, identify faults that are wasting energy or impacting occupant comfort, and then generate performance baselines against which future operation will be compared. When this algorithm is first connected to a new building, an adaptive diagnostic Bayesian network identifies the components and configuration of the AHUs and VAVs, and then generates probabilistic outputs indicating the root causes of potential faults and inefficiencies. Once these issues have been rectified to the satisfaction of the building operator, the algorithm then begins to accumulate training data with detailed information about the system operation (while simultaneously continuing to monitor for additional faults). As this training data is accumulated, the diagnostic confidence of the Bayesian network is continuously improved. Additionally, this use of an operational baseline allows for accurate detection and diagnosis of faults causing gradual performance degradation in addition to faults that abruptly occur.

1. INTRODUCTION

It is well known that the energy costs of operating commercial buildings are high. Commercial buildings represent nearly one-fifth (18.9%) of all U.S. energy consumption, of which 43% is estimated to be spent on HVAC for space conditioning (DOE, 2010). In office buildings specifically, Pérez-Lombard et al. (2008) provided estimates that HVAC energy comprises 48%-55% of the total energy used and that HVAC consumption across all sectors accounts for approximately one fifth of all energy use in developed countries. While HVAC will inevitably comprise a substantial portion of our energy consumption, a portion of this energy is actually being squandered as a result of faulty operation of these systems.

The percentage of total energy wasted as a result of undiagnosed faults is conservatively estimated in the range of 15%-30% of total HVAC energy. Studies have estimated 20-30% energy savings can be achieved by retro-commissioning (Annex 34, 2006) and the most comprehensive longitudinal case study available found a 44% reduction in the use of electricity and a 78% reduction in natural gas consumption across the 10 year study period

(Annex 47, 2010). These data clearly demonstrate that cost-effective and accurate automated fault detection and diagnosis (AFDD) can significantly mitigate the energy wasted in buildings. This study focuses exclusively on AFDD for AHU-VAV systems, but many of the concepts investigated here are applicable to a wider range of systems. In the United States, AHU-VAV systems serve over 30% of commercial floor space (EIA, 2003; EIA, 2013), and they are a known source of many common faults.

This study investigates the use of a Bayesian network for fault isolation in an automated-commissioning (ACx) environment. ACx combines and automates both retro-commissioning (RCx), and continuous-commissioning (CCx) practices. As their names reflect, RCx refers to the identification of existing faults in a building, while CCx refers to continuous monitoring of a building to identify new faults as they occur in real time. These two practices are noted separately because CCx has the added benefit of being able to automatically customize baseline thresholds for each building once RCx has identified any existing faults.

Bayesian networks have been successfully employed for decision support and diagnostics applications in many fields and have recently been included in a number of publications in the HVAC AFDD field for RCx applications. Specifically, the utility of Bayesian networks for AHU-VAV systems were most recently demonstrated for VAV-boxes in Xiao et al. (2014). This work followed up on the paper using the same methodology for AHU fault detection and chiller fault detection (Zhao et al., 2013). This paper utilizes a similar Bayesian network approach, but expands upon the existing approaches to suggest a method for including automated threshold customization to improve diagnostic accuracy in a CCx setting. Additionally, this paper suggests that monitoring multiple levels of severity in the observed evidence in order to improve overall diagnostic accuracy, as well as the inclusion of multiple levels of severity in the fault node in order to provide more accurate and actionable information to a building operator. Beyond the modifications to the network structure, this paper also presents a novel approach for the network parameter estimation. In all previous studies, the network was trained and tested utilizing a single building, but it was found that the networks utilized in these studies suffered when tested on other buildings. During the course of this study, a formal method for identifying generalizable parameters for the Bayesian network conditional probability tables (CPTs) was developed and tested in order to avoid the problem of over-fitting.

2. METHODS

The first step in this study was to generate the overall AHU Bayesian network structure. This was performed based upon review of the literature and the first-principles applicable to an AHU. The Bayesian network is structured such that *fault nodes* represent the different pieces of equipment that can fail, while *evidence nodes* represent the symptoms, or the effects of faults that are observable/measurable in typical building automation system (BAS) data. These fault nodes are then connected via *arcs* to the evidence nodes with which there exists a causal relationship.

2.1 Fault Nodes

Each fault node corresponds to an individual component within an AHU that is being monitored for faults. These components primarily consist of dampers, valves, fans, and sensors. Each node can take on a variety of different values, one of which is always “fault free”. For example, the fault node for an outdoor air (OA) damper can take the following basic values:

- Fault-free (operating normally)
- Stuck too far open
- Stuck too far closed

In order to improve diagnostic accuracy, the “too far open” and “too far closed” were then split to include *moderate* and *extreme* levels of “faultiness”. This is due to the fact that the severity of most faults can have different impacts on the observable evidence as the effects propagate to downstream components. Without this modification, the network would either omit useful diagnostic information, or erroneously identify corroborative evidence as contrary evidence.

2.2 Evidence Nodes

Each evidence node corresponds to a diagnostic *rule* that is based on the operational mode of the system. These rules are typically comprised of a relationship between two or more measured BAS values that can be predicted based upon the operational mode of the system. For example, if the system is operating at 100% outdoor air, it is expected that the mixed air temperature and the outdoor air temperature are approximately equal. Originally these nodes were constructed with three possible values: *higher than expected*, *within normal operational range*, and *lower than expected*. During initial testing it was noted that, while discretizing the evidence nodes was required for computational tractability, the simplification in this manner was perhaps overly-simplified. For example, a temperature deviation of 3°F is much "weaker" evidence than a temperature deviation of 10°F, yet this differentiation was not reflected in the Bayesian network conditional probability tables. As a result, the evidence was binned more finely to include this additional information, so a new node may have values such as *within normal operational range, moderately higher than expected,* and *significantly higher than expected.*

Illustrated in **Figure 1** is a section of the Bayesian network that is utilized to diagnose faults related to the outdoor air damper (OA damper) and mixed air temperature (MAT) sensor. Faults with either of these items can cause significant energy waste, negatively impact indoor air quality, and can also affect occupant comfort during economizing operation. Additionally, faults with these two components are known to go undetected for long periods of time, and have similar fault "evidence", making it potentially difficult to rapidly identify the root cause of the fault. The term "fault evidence" is used interchangeably with the term "fault symptom" and refers to the behavior of a fault that can be observed in the BAS data.

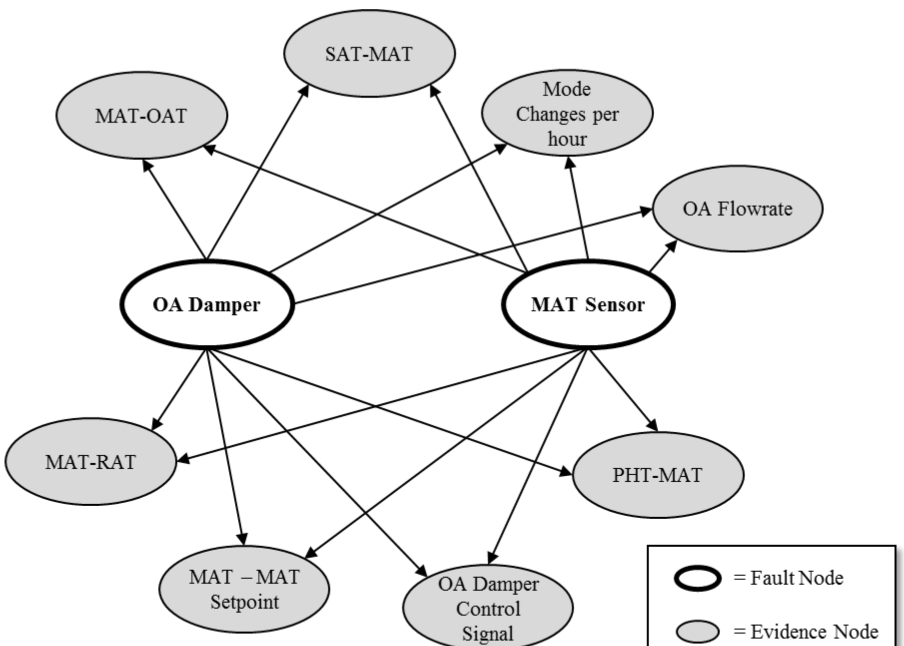

Figure 1: Bayesian network excerpt diagram (economizer section)

Figure 1 utilizes the following abbreviations: OAT for *outdoor air temperature*, MAT for *mixed air temperature*, SAT for *supply air temperature,* RAT for *return air temperature,* PHT for *preheat (between-coil) temperature*.

2.3 Bayesian Network Parameter Estimation

The models used in this study are considered graded, leaky, causal, noisy-max models. (Depending on which reference is utilized, there is some variation in nomenclature, but they will be referred to simply as noisy-max models throughout this document.) The "noisy" aspect takes into account that the fault may not always cause the expected effect in the evidence. The "leaky" aspect refers to the fact that there is some probability that a faulty evidence state will be observed without any underlying fault. The causal aspect refers to the fact that all of the arcs connecting the nodes must be utilized only for direct cause-effect relationships. And finally, the graded aspect refers to the fact that the noisy-max model can only be used for graded variables, since it is a max function. Good explanations of the practical applications of noisy-max models can be found in Zagorecki & Druzdel (2013). Use of the noisy-max model provides two key advantages in the use of generating Bayesian networks: (1) a reduction in the number of probability states required to be specified to generate the conditional probability tables, and (2) the ability to be reduced for more efficient computation (Díez & Galán, 2003). This use of the canonical model significantly simplifies the process of selecting the CPT parameters and standardizing these parameters for multiple buildings. Previously, use of Bayesian networks for HVAC diagnostics required customized for a specific building, but in order for this technique to be viable for commercial applications it must be applicable to all AHU-VAV systems.

The first step was to significantly reduce the number of possible options that the network designer would have to choose from for each combination of fault and evidence. For a typical evidence node, if random probability values were selected for analysis, there could be nearly two million different potential combinations to evaluate. Evaluating the efficacy of these values is very time consuming for each potential combination, so performing a traditional optimization is not feasible. As a result, the probabilities are binned, and the values are based upon the known properties of AHU-VAV systems. To perform this simplification, a set of bins were defined as follows:

- The evidence state is the definition of the fault state in question
- The evidence state is likely to be observed, given the fault state in question
- The evidence state has a "split-likelihood", given the fault state in question. For example, in an instability fault, the evidence is equally likely to be too high as it is to be too low.
- The evidence state is unlikely to be observed, given the fault state in question – but there exists some causal link between the fault and evidences states in question.
- There is no causal link between the fault and evidence states in question.

Since these definitions are well understood for most fault/evidence combinations, assigning each arc to one of these combinations is possible based on first-principles. There may be some advantage to adding some other options beyond the five described above, and that can be investigated in the future. Once these states are assigned, the next step was to provide a standard value for each possible definition. The default values selected are included in **Table 1**, below.

Table 1: Bayesian network parameters

Abbrev	**Description**	**Prob. Value**
3 Evidence States (noisy-max)		
D	Evidence is definition of fault, or fault-free state	1
L_f	Likely, extreme severity	0.5
L_m	Likely, moderate severity	0.5
c_2	Likely remainder	= 1 - L_f - L_m
S_f	Split likelihood, extreme severity	0.25
S_m	Split likelihood, moderate severity	0.25
c_3	Split likelihood remainder	= 1 - S_f - S_m
U_f	Unlikely, but possible, extreme	0.05
U_m	Unlikely, but possible, moderate severity	0.05
c_1	Unlikely, but possible, remainder severity	= 1 - U_f - U_m
N	Not related at all	0
2 Evidence States (noisy-or)		
L_2	Likely (2 states)	= L_f + L_m
S_2	Split likelihood (2 states)	= S_f + S_m
U_2	Unlikely (2 states)	= U_f + U_m

In this table it can be observed that only six parameters (highlighted in black) are eligible to be adjusted in order to perform refinement of the method. Since the analysis of the network across many different fault scenarios and buildings is computationally- and time-intensive, limiting the number of parameters allows for a future improvement via sensitivity analyses and informal optimization approaches. Since no tuning of these parameters was performed using real building data – only fundamental knowledge of general AHU-VAV variable interactions – the model is not over-fit for any specific building, but generally constructed to apply to all buildings. Any future improvements or refinements should be performed utilizing testing across multiple buildings to ensure that this remains the case.

2.5 Field Testing

Field testing of this approach was performed at three occupied commercial buildings in the Philadelphia area with AHU-VAV systems. In each of these buildings a variety of faults were artificially introduced to mimic commonly observed faulty conditions. For different fault scenarios, this was performed either via control system override, physical manipulation of components, or by manually applying a voltage control signal to a system component. The method utilized for each fault experiment was dictated by the ability to implement each fault in a manner that would accurately mimic not only the physical effects of the fault, but also accurately mimic what the BAS data would be if the fault had been naturally occurring.

Building 1

The first building in which fault experiments were conducted was Building 101 at the Philadelphia Naval Yard. It is a three-story 55,000 ft^2 commercial building located on the Navy Yard campus at the southern edge of Philadelphia at the confluence of the Delaware and Schuylkill rivers. The Building 101 HVAC system consists of three AHUs serving 26 VAV boxes. The heating is supplied by a natural gas boiler and the cooling is provided by three direct expansion units. The HVAC system is controlled by a BACnet system, with a typical AHU sensor arrangement. This building has a standard AHU-VAV configuration, except the AHUs do not have mixing boxes – instead having a small outdoor air damper with no return or mixed air dampers. The outdoor air flowrate is minimal, so the building ventilation is primarily due to infiltration through the envelope.

Building 2

Building 2 is the Swope Music Building, a three-story, 90,000 ft^2 school building at West Chester University. Located approximately 23 miles west of Philadelphia, this building was constructed in 2007 and received a silver LEED certification at that time. This building contains a 375-seat theater, a 125-seat recital hall, a music library, and numerous rehearsal spaces, classrooms, and offices. The HVAC system is served by 7 AHUs, with space-specific AHUs for the Theater Lobby, the Theater, the Instrument Rehearsal Room, the Recital Hall, and the Stage. Two other AHUs serve the 90 VAV-boxes that condition the spaces throughout the rest of the building. In addition to the typical AHU-VAV configuration, this building requires careful humidity control so there are additional humidity sensors and controls utilized with this system.

Building 3

Building 3 is Stratton Hall at Drexel University, a three-story 50,000 ft^2 building constructed in 1955, with a complete renovation of the HVAC system and improvements to the building envelope performed in 2013. This building has a typical AHU-VAV system, with separate AHUs serving the

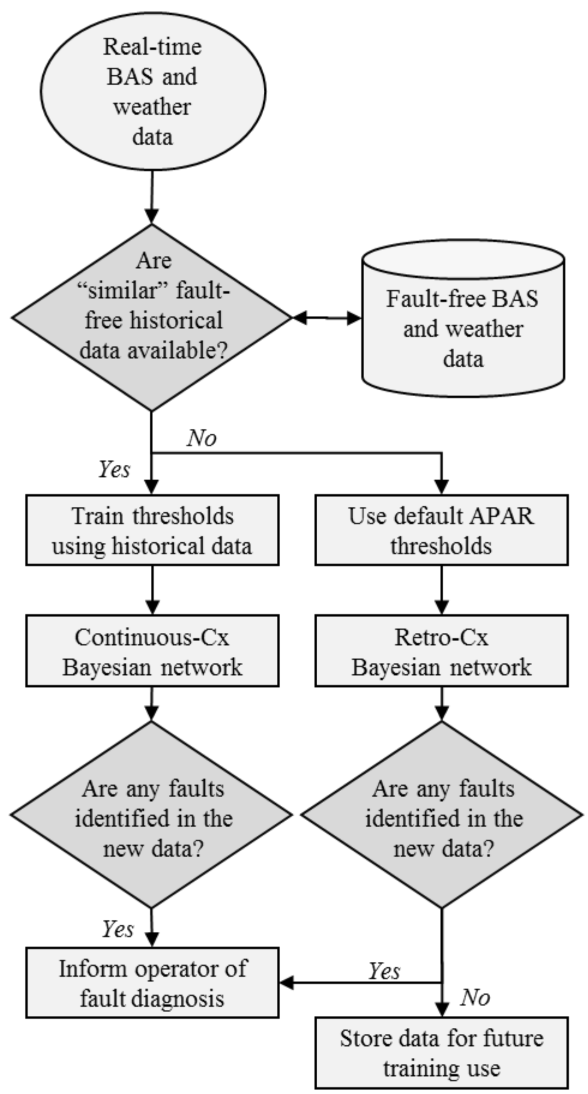

Figure 2: RCx versus CCx scenarios

second and third floors of the building. These AHUs are served by a district chilled water plant and a heat exchanger connected to Philadelphia's steam heating infrastructure. The first floor of this building utilizes small window units and direct expansion units for cooling and radiant heat, that are not connected to the BAS, so only the second and third floors were included in this study.

Fault Experiments

To perform initial testing of this method, 15 fault experiments were selected from each building for analysis. Each building was studied in two ways: simulating the RCx scenario, and simulating the CCx scenario. Due to space constraints, a comprehensive listing of all faults studied is not included, but the faults were selected to obtain a cross-section of different fault "types" and severities. The faults selected included fan faults, coil/valve faults, damper faults, and sensor faults (temperature and pressure sensors) at each building. These faults were implemented across a variety of severities, with some faults causing noticeable impact to the system operation and other faults causing negligible impact to overall system operation.

In a commercial deployment setting, the RCx scenario would be the first to be tested, and then once the data is identified as a suitable baseline, this information would be used to train the CCx network, as depicted in **Figure 2**. In this figure, the term "similar fault free historical data" refers to data from the same mode of operation as is being analyzed that has already been found to be free of significant faults using the RCx network.

3. RESULTS

To analyze the results, the day during which each fault was implemented was analyzed, and the results of the Bayesian network were summarized over the duration of each fault experiment. If a fault was identified with an average probability exceeding 60% during the time in which it was implemented, it was considered to be *accurately diagnosed*. If a different fault was diagnosed with higher than 60% confidence, and this incorrect diagnosis was in-part due to the effects of the fault that was artificially introduced, it was categorized as *mis-diagnosed*. If a different fault was identified during the course of the fault experiments that was not due to the fault experiment, it was categorized as a *false-alarm*. If no fault was detected with greater than 60% confidence, the fault was categorized as *not detected*. The faults that were not detected were also sub-classified as either *impactful* faults of faults with *no impact.* A fault is only categorized as having no impact if:

- It is not in the control-loop of the system, so it does not directly impact any operation
- If an experienced building operator was reviewing at the building data, with knowledge that the fault exists, it would not be apparent from the BAS data.

The first bullet above is objectively determined, but the second bullet requires a subjective assessment of whether the fault has any system impact. The faults that were found to have no impact were sensors with very small biases (2-3°F), and dampers or valves that were "stuck" at their normal operational range for a given operational mode. For example, if an outdoor air damper control signal is requesting the damper to be 45% open, but the fault experiment has the value fixed at 50% open, the difference in the data would be very difficult to discern and the energy, cost, and comfort impacts would be negligible. However, when the system shifts to a different mode and requests that the damper is closed, or moved to 100% open, then the fault would become immediately apparent to a building operator, and would have measurable impact on the system operation.

A summary of the results of the faults utilizing the RCx approach (baseline threshold values) is included in **Table 2**, and a summary of results utilizing the CCx approach (training the Bayesian network thresholds based on known fault-free operation) is included in **Table 3**.

Table 2: Retro-commissioning summary

Retro-Commissioning Results						
Building	Total Faults	Accurately Diagnosed	Misdiagnosed	False Alarms	Not Detected	
					Impactful	No Impact
Building 1	15	10	1	1	0	4
Building 2	15	9	2	2	0	4
Building 3	15	12	0	0	0	3

Table 3: Continuous-commissioning summary

Continuous-Commissioning Results						
Building	Total Faults	Accurately Diagnosed	Misdiagnosed	False Alarms	Not Detected	
					Impactful	No Impact
Building 1	15	11	0	0	0	4
Building 2	15	11	0	0	0	4
Building 3	15	13	0	0	0	2

From these tables, it can be firstly observed that no *impactful* faults were missed in either network, and that when using the continuous commissioning network, all faults that impacted energy, comfort or indoor air quality were accurately detected and diagnosed. To illustrate the concept of a fault that is not impactful, the operational data for the fault in which the Building 2 return fan was stuck at a fixed speed is shown in **Figure 3** and the diagnostic result is included in Figure 4. As observed in **Figure 3**, the fan speed was fixed at the normal operational range, resulting in negligible impact to the system, and limited evidence to use for diagnosis.

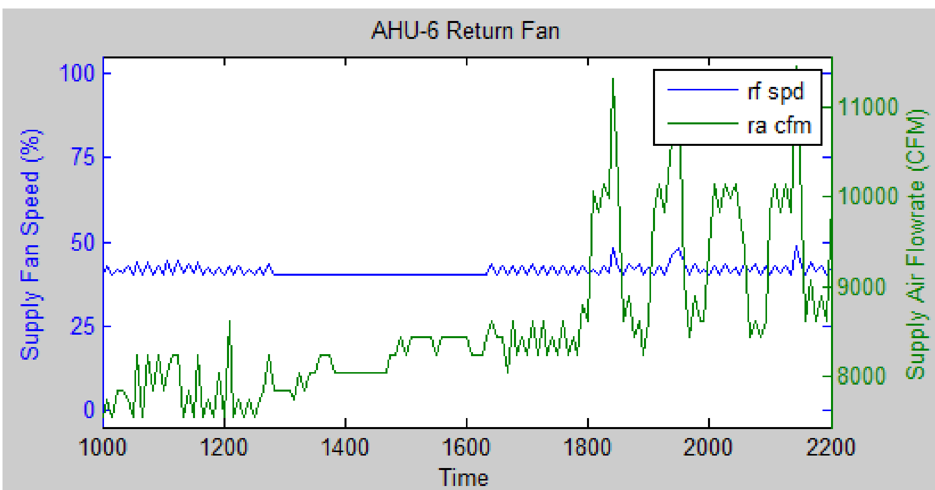

Figure 3: Return fan stuck at fixed speed (minimal impact)

In the above figure, “rf spd” refers to the return fan speed, and “ra cfm” refers to return air flowrate.

As demonstrated in **Figure 4**, below, the most likely diagnosis was correctly found to be a fan stuck at fixed speed ("frozen"), but due to the limited impact the diagnosis was not strong enough to meet the 60% threshold used to evaluate the efficacy of the fault diagnostic algorithm.

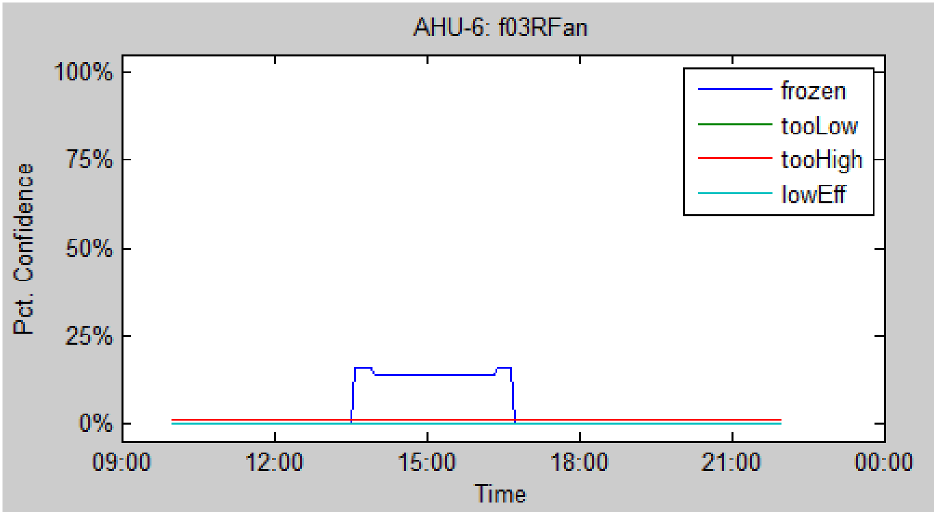

Figure 4: Return fan diagnostic result, 2015-12-17 (Swope AHU-6)

In the retro-commissioning network, there were three false alarms and three misdiagnosed faults. In reviewing the results, it was found that the misdiagnoses corresponded directly with the false alarms. The false alarm and misdiagnosis in Building 1 were due to the fact that there is no return air (RA) or mixed air (MA) damper. As a result, when economizing is expected, little to no outdoor air is actually being introduced into the system. While this was not technically a fault (since no piece of equipment was malfunctioning), it is behavior that differs from typical AHU operation, so it is useful information for a building operator or service provider. However, when previous operational data was used to train the thresholds, these faults were not falsely flagged.

Similarly, the misdiagnoses and false alarms identified in Building 2 were due to the non-traditional humidity-based control that was utilized in this building. In one instance a fault was flagged due to simultaneous operation of the heating and cooling coils. While this was done intentionally at this site for the purpose of dehumidification, this is indicative of faulty operation in many buildings and is a reasonable operational condition of which to alert a building operator or service provider. The other false alarm, and associated misdiagnosis was due to the temperature differential created by the steam humidification. Again, once these items were incorporated in the training data as "normal" operation, the faults were properly diagnosed.

4. DISCUSSION

This demonstration of the AHU-VAV AFDD tools provides evidence that the use of a Bayesian network approach can be an effective diagnostic tool and can be utilized across multiple AHUs to diagnose all different categories of faults in both RCx and CCx applications. The use of a "passive" diagnostic algorithm (not manipulating system controls) will always be limited by the impact of the fault on system operation, but the method utilized by this AFDD tool allows for accurate detection and diagnosis of faults across a vast majority of AHU-VAV configurations and applications.

The results presented here are the initial results from the Bayesian network investigation, prior to any tuning or refinement being performed. The network parameters for this initial study were selected based upon expert knowledge and first principles. An effort is underway to refine the parameters using the three existing "training" buildings, and tested on additional buildings to see if the results are improved. These initial results are very promising, but further testing and validation of this method is required in order to properly quantify the strengths and weaknesses. This can only be accomplished through extensive additional testing across multiple buildings.

5. CONCLUSIONS

The primary conclusions from this study include:

- The use of a Bayesian network is an effective method for combining multiple rules in order to identify the root cause (source) of abnormal/faulty/sub-optimal AHU operation.
- By generalizing the parameter selection based on first-principles of a typical AHU system, it is possible to create a diagnostic network that is effective across multiple buildings with very limited tuning.
- Allowing the Bayesian network to automatically customize thresholds for individual buildings fault-free baselines can improve accuracy over a general Bayesian network.
- There appears to be some advantage to including different levels of severity for fault and evidence nodes in terms of diagnostic accuracy. Further work needs to be done to quantify this improvement.

REFERENCES

ANNEX 25, (1996) *Building Optimization and Fault Diagnosis Source Book, ANNEX 25 Real Time Simulation of HVAC Systems for Building Optimization, Fault Detection and Diagnosis*, International Energy Agency, Finland.

Díez, F. J. & Galán, S.F. (2003). Efficient computation for the noisy max. *International Journal of Intelligent Systems*, *18*(2), 165-177.

Katipamula, S., & Brambley, M. R. (2005a). Methods for fault detection, diagnostics, and prognostics for building systems - A review, Part I. *HVAC and R Research, 11*. 3-25.

Katipamula, S., & Brambley, M. R. (2005b). Methods for fault detection, diagnostics, and prognostics for building systems - A review, Part II. *HVAC and R Research, 11*. 169-187.

Pérez-Lombard, L., J. Ortiz, et al. (2008). A review on buildings energy consumption information. *Energy and Buildings 40*(3). 394-398.

Xiao, F., Zhao, Y., Wen, J., & Wang, S. (2014). Bayesian network based FDD strategy for variable air volume terminals. *Automation in Construction*, *41*, 106-118.

Zhao, Y., Xiao, F., & Wang, S. (2013). An intelligent chiller fault detection and diagnosis methodology using Bayesian belief network. *Energy and Buildings, 57*, 278-288.

Zagorecki, A. & Druzdzel, M. J. (2013). Knowledge Engineering for Bayesian Networks: How Common Are Noisy-MAX Distributions in Practice?. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, *43*(1), 186-195.