Title

Detecting Passing Valves at Scale Across Different Buildings and Systems: A Brick Enabled and Mortar Tested Application

Authors

Duarte Roa, Carlos
Raftery, Paul
Singla, Rupam
et al.

Publication Date

2022-08-01

DOI

10.20357/B7VP5H

Building Technologies & Urban Systems Division
Energy Technologies Area
Lawrence Berkeley National Laboratory

Detecting Passing Valves at Scale Across Different Buildings and Systems: A Brick Enabled and Mortar Tested Application

Carlos Duarte Roa[1], Paul Raftery[1], Rupam Singla[2], Marco Pritoni[3], Therese Peffer[4]

[1]Center for the Built Environment, University of California
[2]TRC
[3]Lawrence Berkeley National Laboratory
[4]California Institute for Energy and Environment

Energy Technologies Area
August 2022

Duarte Roa C., Raftery P., Rupam S., Pritoni M., Peffer T. (2022). Detecting Passing Valves at Scale Across Different Buildings and Systems: A Brick Enabled and Mortar Tested Application. ACEEE Summer Study on Energy Efficiency in Buildings 2022 https://doi.org/10.20357/B7VP5H

This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy,
Building Technologies Office, of the US Department of Energy
under Contract No. DE-AC02-05CH11231.

Detecting Passing Valves at Scale Across Different Buildings and Systems: A Brick Enabled and Mortar Tested Application

Carlos Duarte Roa[1], Paul Raftery[1], Rupam Singla[2], Marco Pritoni[3], Therese Peffer[4]
[1]Center for the Built Environment, University of California; [2]TRC; [3]Lawrence Berkeley National Laboratory; [4]California Institute for Energy and Environment

KEYWORDS

Hot water distribution, Variable air volume, Passing valve, Fault detection and diagnostics, Energy efficiency

ABSTRACT

Heating hot water distribution systems are typically used in commercial buildings to condition spaces to provide occupant thermal comfort. However, recent research shows significant distribution losses within these systems that drive down the overall hot water plant efficiency. This research focuses on detecting passing valves in reheat coils found in variable air volume (VAV) terminal units to reduce distribution losses. A passing valve allows hot water flow when the actuator on the valve is commanded to be closed. The fluid causes unintentional heating or cooling to occur, causing comfort and control issues, and wasting energy. We developed the passing valve detection algorithm using a framework based on the Brick schema and Mortar platform to ensure that the application is portable and can scale to many buildings. We applied the same application to analyze 1,335 VAV reheat terminal units in 20 buildings. The diversity found in these large datasets increases confidence that any building with VAV reheat terminal units with the required sensors and Brick data model can run our open-source algorithm with little or no modification. In aggregate, 5% of VAV units analyzed were categorized as having a sensor fault, 14% with potential passing valve fault, and 81% with no faults detected. However, there is a significant variation in the proportion of VAV units with a passing valve detected (1% to 83%) of each building's analyzed units.

Introduction

Hot water distribution systems with natural gas-fired boilers as the heating energy source are widely used in heating, ventilation, and air-conditioning (HVAC) systems for large commercial buildings (EIA 2022). The hot water distribution system makes use of valves to control the flow of hot water to centralized (e.g. air handling units) and distributed (e.g. variable air volume (VAV) terminal units) HVAC equipment. However, recent research has shown that hydronic heating systems may have substantial distribution losses and other inefficiencies that are not accounted for in current practice (Raftery et al. 2018). Several pathways exist where distribution losses can occur but, in this research, we focus on identifying passing valves in VAV reheat terminal units using open-source software that can scale to a large number of heterogenous buildings with little or no modifications to the developed software.

A passing valve allows hot water to flow when it is expected to be closed. It occurs when the valve seal has failed due to repeated stress forces, or due to an old and brittle seal, or due to fouling or a blockage. It can also occur in an automated valve when the valve is stuck due to fouling, a failed motor, faulty internal equipment, or loss of communication with the building management system. In some cases, a valve may also be manually overridden to a fixed position for the purposes of testing or to resolve a temporary comfort issue, which is then forgotten and left that way. Passing valve failures can be long-term as in the example of a failed valve seal or short-term when valve sticks on occasion or intermittently loses communication with its control system. Passing valves in HVAC systems are a common fault that goes unnoticed due to feedback masking and compensatory action of other downstream equipment (Salsbury and Diamond 2001; Najafi 2010). Identifying passing valves may be labor intensive where facility personnel must inspect valves visually or by using specialized tools to perform valve diagnostics using acoustical measurements (Kaewwaewnoi, Prateepasen, and Kaewtrakulpong 2010), infrared thermography (Balaras and Argiriou 2002; Zhao et al. 2021), or manual review of the time-series data that is acquired the existing building automation system sensors. The alternative for passing valve detection is to use a data-driven approach.

The operation of buildings increasingly relies on a network of sensors and Internet-of-Things (IoT) devices that produce unprecedented amounts of data. However, these data typically go unused beyond the operation of the building. Consequently, there are missed opportunities from underutilizing this resource such as energy consumption, cost, and control optimizations, predictive maintenance, visualization and reporting, and fault detection and diagnostics. These missed opportunities may cost a building 5%-40% in energy savings (Ahmed et al. 2010; Lin, Kramer, and Granderson 2020). Some of the barriers identified that hinder the widespread use of advanced analytics and controls include proprietary equipment and building management systems, unique naming convention of building assets and data points, and the inherent uniqueness of buildings and their systems (Ahmed et al. 2010; Fierro 2021). All these barriers contribute to the lack of interoperability and portability of software tools that enable the use of all these advanced features as standard practice. In other words, it is typical that advanced analytic implementations for buildings become unique deployments specific to a building.

Application development

Background

We developed our passing valve detection application using a framework based on Brick and Mortar to ensure that the application is portable and can be applied to many buildings. The prerequisites to run the application are a Brick data model of a portion of a building containing VAV reheat terminal units, and data from at least three data streams. The required sensors are an upstream and downstream air temperature sensor of the reheat coil. The other data stream is the valve position from either a sensor or a command. The upstream temperature sensor here is typically the supply air temperature sensor on the air handling unit (AHU). The downstream temperature is the leaving air temperature from the VAV unit also known as the VAV supply or discharge air temperature. The relationship between which AHU in a building serves which VAV box is automatically extracted from the Brick data model of that building.

A Brick model is a semantic metadata model structured using the Brick ontology which provides standardized descriptions of the physical, logical, and virtual assets found in buildings and formal definitions to establish relationships between the assets (Balaji et al. 2016). The Brick schema organizes classes into a hierarchy with three root classes: *Equipment*, *Point*, and *Location*. The definitions of subclasses from these root classes become more specific further down the class hierarchy. For example, *Equipment* has subclasses *HVAC*, *Lighting*, *Electrical* while each of these has its own subclasses such as *Air Handling Unit* and *Boiler* for HVAC and so on. Eventually, we can identify a class in the hierarchy that is specific enough to assign to a building's asset and associate properties to it. Then we can take two classes and associate them by assigning a formal Brick relationship between them. Figure 1 shows how we use the relationships *hasPoint*, *hasPart*, and *feeds* for a Brick data model of a VAV with reheat terminal unit.

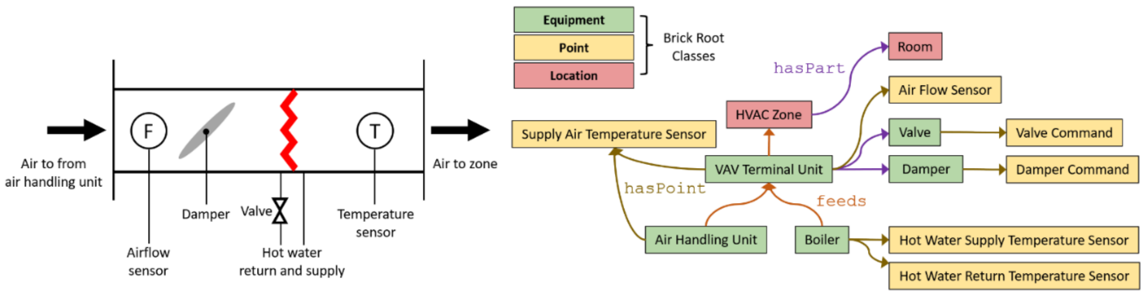

Figure 1: Left) schematic of a variable air volume (VAV) with reheat terminal unit and right) a Brick data model of the VAV terminal unit.

New algorithms that improve building performance are constantly being developed but are inconsistently evaluated, and it is often the case that an algorithm is tested using only one building (Miller, Nagy, and Schlueter 2018; Miller 2019). Thus, there is a lack of understanding of an algorithm's generalizability or how it might perform on different buildings. Mortar addresses this issue by providing an open testbed and platform for developing and evaluating algorithms for the built environment (Fierro et al. 2020). It contains timeseries data and corresponding Brick models for over 100 buildings. The large number of buildings not only provides variety in the type, location, design, equipment, occupant behavior, and operational strategies found in those buildings but also variety in the approach that various Brick modelers took to build the buildings' data model. These variations in the data models allowed us to generalize programmatic queries that retrieve the necessary sensor data. Brick allows us to retrieve data by its purpose, behavior, and context instead of using the non-standard naming conventions assigned by a particular building management system vendor. Creating a generalized query can be a challenge but research to automatically generalize them is underway (Bennani et al. 2021). Thus, Mortar was a critical component for our application development, where we tested the generalizability and impact of our application.

Procedure to detect passing valves

The following sections describe the cleaning, preprocessing, and analysis procedures to detect passing valves in VAV reheat terminal units using the building management system data. The minimum data streams required for this application to run are the supply air temperature from both the AHU and the VAV terminal unit and the VAV hot water valve position. If an air flow sensor is available for the VAV unit, then the application will use it to increase the reliability for filtering data when the unit is in operation. The air flow rate measurements will reduce false positives in the application's detection of passing valves by letting that application know that air temperature sensors are measuring moving air and not static air that might be at room temperature.

Data cleaning and preprocessing. The application starts by reading the building's Brick data model and applying the predefined programmatic query to identify the required and optional data streams. The data is downloaded from its source if the query can find the required data stream classes within the Brick model. We applied the application to the Mortar database and two other site datasets. Most of the data streams in the datasets are at 15-minute intervals or less and have not undergone a cleaning process. We first ensure that each data stream from a VAV terminal is at the same time interval by downsampling each stream to the largest time frequency (rounded to the nearest five minutes) found from all streams. Then we perform several checks to confirm that all required data streams are available, for each timestamp there is data available from all required data streams, and the sequence of timestamps is continuous for proper analysis. We flag instances when there are missing timestamps in the data which can affect the data analysis in subsequent steps e.g., determination of transient versus steady-state conditions.

Data analysis. The first analysis we performed on each VAV unit dataset was to separate the transient response of the VAV unit from the steady-state response. We located all valve switchovers from open to closed and vice versa. For each open to closed switchover, we assumed steady-steady conditions after 12 minutes of the valve closing. The time threshold is based on Raftery et al. (2018) calculations that showed that, in a typical VAV box application, approximately this time frame is sufficient to dissipate the remaining heat in the hot water coil after the valve closes.

Next, we verify that the sensor measurements are not stuck at a constant value, have reasonable variance, and not at extreme values. Once sensors pass these checks, we use the air flow sensor, if available, to determine when the VAV unit is in operation. Air flow rates can vary widely between individual VAV units based on the zone they supply. In addition, non-ideal air flow inlet conditions can exist where the air flow rate measurement accuracy is affected (Liu, Wen, and Waring 2014). Thus, we cannot use a single air flow rate threshold value. Instead, we use kernel density estimation (KDE) to determine the probability density function of air flow rates for each VAV terminal unit. This process creates a multimodal distribution as shown in Figure 2 where we can infer the no air flow condition (first peak) and typical operation which is most likely at the minimum air flow rate condition (second peak). Knowing the location of the first peak and the first trough gives us a range to select a reasonable air flow rate to consider the unit in operating mode. We arbitrarily selected 40% of the range and added it to the first peak to calculate the minimum air flow rate cutoff. Users may adjust this “accuracy parameter” percentage depending on their certainty of the air flow measurement. If air flow measurements are not available, we use typical building hours (6:00 to 18:00) on weekdays to filter for VAV unit operation data.

Next, we calculated an initial median temperature difference between the upstream and downstream temperatures of the VAV reheat coil when the valve is closed. We only used data that is in steady-state and above the minimum air flow rate cutoff or within typical building hours which we refer to as the *operational* dataset. We then determined a threshold, defined as the minimum of either two times the initial median temperature difference or a predefined value of 10 °F, to filter out probable passing valve operation. That is, if the temperature difference for a closed valve is above the threshold, then we would consider it probable passing valve operation. We used the remaining typical closed valve operation points to update the median temperature difference. We then used a four-parameter sigmoid function to model the relation between temperature difference and valve operation. We assumed that the updated median temperature difference to be the lower plateau and the 95th percentile of the measured temperature differences in the VAV unit to be the upper plateau of the sigmoid function. We used non-linear least squares as implemented in the SciPy Python package optimization modules to find the slope and the center of the sigmoid function (Virtanen et al. 2020).

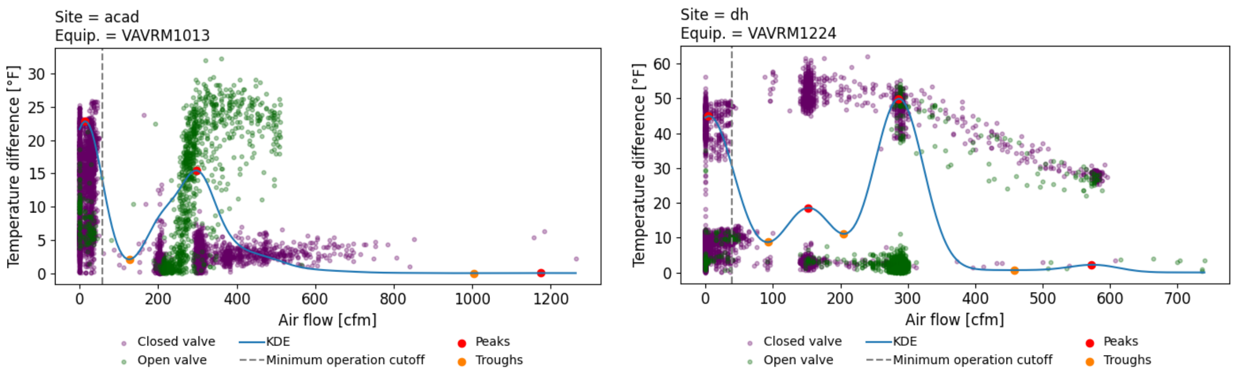

Figure 2: Temperature difference between the upstream and downstream air temperature of VAV reheat coil versus its air flow rate when reheat valve is open (green) and closed (purple). The solid blue line is the kernel density estimation (KDE) with peaks (red) and troughs (orange) highlighted. The dashed gray line indicates the calculated minimum air flow rate to consider the VAV in operating mode.

The model allows us to refine the temperature difference threshold at which we would consider passing valve operation. We define the refined threshold at the point in the model where the bottom plateau transitions into the linear portion of the sigmoid function i.e., threshold located in the knee of the first bend. Thus, if the temperature difference in the *operational* dataset when the valve is closed exceeds this threshold for four consecutive hours, then we consider it a long-term passing valve event. If the threshold is exceeded for one hour, it is a short-term passing valve event. There can be multiple fault events in each VAV unit dataset. We would need to detect at least two long-term passing valve events to categorize the VAV unit as having a passing valve fault or if the total time of fault operation exceeds 5% of the total VAV dataset time period. In the case of short-term events, we would categorize it as passing valve fault operation if the 5% time exceedance threshold is met. We arbitrarily selected 5% as a reasonable starting percentage but more research is needed to evaluate the optimal value on this parameter and others used in the application. As building operators typically have limited resources for evaluating and addressing faults, the most viable approach would be to rank the faults in terms of both probability of a true positive (i.e., a ‘real’ passing valve), the amount of time the fault occurs in the dataset, and the magnitude of the fault (a high temperature difference on a larger airflow matters more). With this information, the building operators can easily prioritize and address the most impactful faults first. For facilities teams with multiple buildings, the fault data could also be prioritized based on aggregate faults at the whole building level, which would allow them to address the most impactful buildings first, instead of zones, and may be a more efficient use of resources.

Results

We analyzed three datasets that contain a total of 20 buildings and 1,335 VAV reheat terminal units. Most of the data come from buildings in the Mortar database and the total number of VAV units only includes units with the minimum amount of data streams for analysis. We selected the timeframe for the Mortar database to range from January through June 2018, October through November 2021 for the ‘bear’ dataset, and May through November 2021 for the ‘lion’ dataset. Table 1 contains summary statistics for each building analyzed using the detection algorithm described above. In aggregate, the algorithm categorized 5% of VAV units as having a sensor fault, 14% with a potential passing valve fault, and 81% with no faults detected. However, when looking at individual buildings with 25 VAV units analyzed or more, the variation in passing valve faults detected has a large range from 1% to 83% of the total building’s units analyzed. Furthermore, the median temperature rise with a closed valve as air is distributed from the AHU’s supply air temperature sensor to the VAV sensor is 1.8 °F for datasets with no faults detected and 6.1 °F for datasets where passing valve faults were detected.

Table 1: Summary statistic of the analysis of VAV reheat terminal unit data streams.

Site[1]	VAV count	Sensor fault	Valve fault	Temperature rise[2] [° F] Median (IQR)		Heat loss due to passing valve[3]
				No Fault	Valve Fault	
brig	418	-	7%	1.6 (0.9-3.1)	13.6 (7.1-27.8)	-
hutch	134	-	4%	1.7 (1.3-2.4)	4.3 (2.9-8.2)	-
chem	117	2%	6%	2.4 (0.6-3.1)	8.8 (6.3-9.5)	-
lion[1]	109	-	10%	2.1 (1.4-2.8)	3.0 (2.5-4.1)	4%
stor	102	-	1%	1.1 (0.7-1.7)	3.2 (3.2-3.2)	-
acad	87	-	15%	2.4 (1.1-6.7)	8.8 (8.2-9.8)	10%
eps	70	-	9%	3.0 (1.3-3.7)	4.3 (3.6-4.6)	93%
dh	57	-	42%	1.6 (1.2-2.0)	2.5 (1.9-2.9)	5%
gha_ics	52	-	83%	4.0 (2.9-5.4)	4.9 (4.4-6.1)	47%
arc	50	76%	20%	8.8 (8.7-8.8)	8.5 (8.2-8.9)	271%
vm3a	35	6%	51%	1.7 (1.4-2.5)	8.0 (4.7-17.1)	554%
bear[1]	29	-	7%	0.9 (0.2-1.5)	3.4 (3.3-3.4)	2%
crus	24	-	-	2.8 (1.4-3.6)	-	-
wsrc	14	64%	36%	-	7.4 (7.4-7.4)	-
mann	8	-	63%	3.5 (2.1-4.4)	6.5 (6.3-9.7)	-
artx	8	100%	-	-	-	-
giedt	7	86%	14%	-	16.3 (16.3-16.3)	-
bwfp	7	-	14%	1.4 (1.2-1.7)	13.3 (13.3-13.3)	-
junger	5	-	-	1.0 (0.3-1.1)	-	-
music	2	-	50%	3.0 (3.0-3.0)	9.3 (9.3-9.3)	-
All Datasets	**1335**	**5%**	**14%**	**1.8 (1.0-3.0)**	**6.1 (3.7-9.3)**	**8%**

1. All sites except bear and lion come from the Mortar dataset.
2. Median and interquartile range (IQR) of temperature rise between upstream and downstream air temperature sensors when the VAV reheat valve is commanded close.
3. Heat loss is calculated for VAV terminal units where a passing valve was detected and all required data streams and air flow rate are available. The heat loss is presented as a fraction of total intentional reheat energy used by the VAV during the analysis period.

The median long-term passing valve fault events detected were 6 and 10 for short-term passing valve fault events per VAV units where either event was detected. The median duration of these fault events is 355 and 128 minutes for long- and short-term events, respectively. These faults resulted in an average heat rate loss of 1,375 Btu/hr with a cumulative 14,400 kBtu of energy loss or 8% of the total intentional reheat energy used by all VAV units analyzed for the time period. We calculated the heat rate loss according to the modified heat balance equation found in Raftery et al. (2018)**.**

Figure 3 shows *operational* data examples from VAV units with various faults illustrated. It includes passing valve, sensor, and reversed valve faults. The dotted light blue line in Figures 3 and 4 represents the revised median temperature difference when the valve is closed, dashed light red line represents the median temperature difference when passing valve fault operation (red dots) is detected with a closed valve, and dashed purple line represents the model developed for the specific VAV unit. The fault operation proportion is the proportion of passing valve fault time (red dots) to the total time period of the individual VAV unit dataset. A reversed valve fault is when the wiring or the control logic is reversed such that the command signal would close the valve instead of opening it. A closer inspection of the data with the sensor fault shown in Figure 3 shows that the supply air sensor of the VAV unit is reporting a constant value of 74 °F.

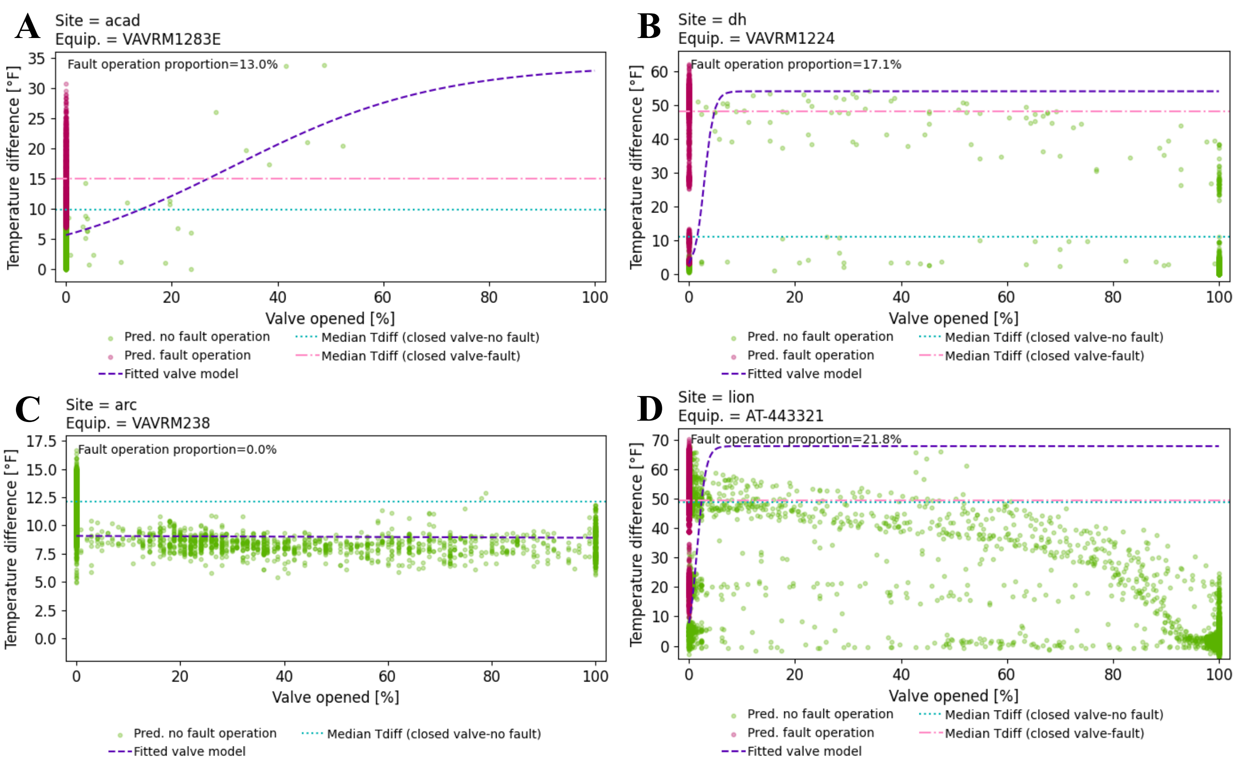

Figure 3: *Operational* data examples from VAV units with various faults. A), B) Passing valve fault, C) sensor fault, and D) reversed valve fault. A high median temperature difference when the valve is closed (red dots) is an indication for a potential passing valve.

Figure 4 shows *operational* data examples from VAV units where no faults were detected. It also illustrates when the air flow rate and typical building hours were used to filter for the operation mode of the VAV unit. Using typical building hours to filter for operation mode seemed to work well. The typical building hours seem to capture a wide range of heating coil responses as the valve command varied. For VAV units with an air flow sensor available, the median minimum air flow rate threshold to consider a VAV in operating mode was 72 cubic feet per minute.

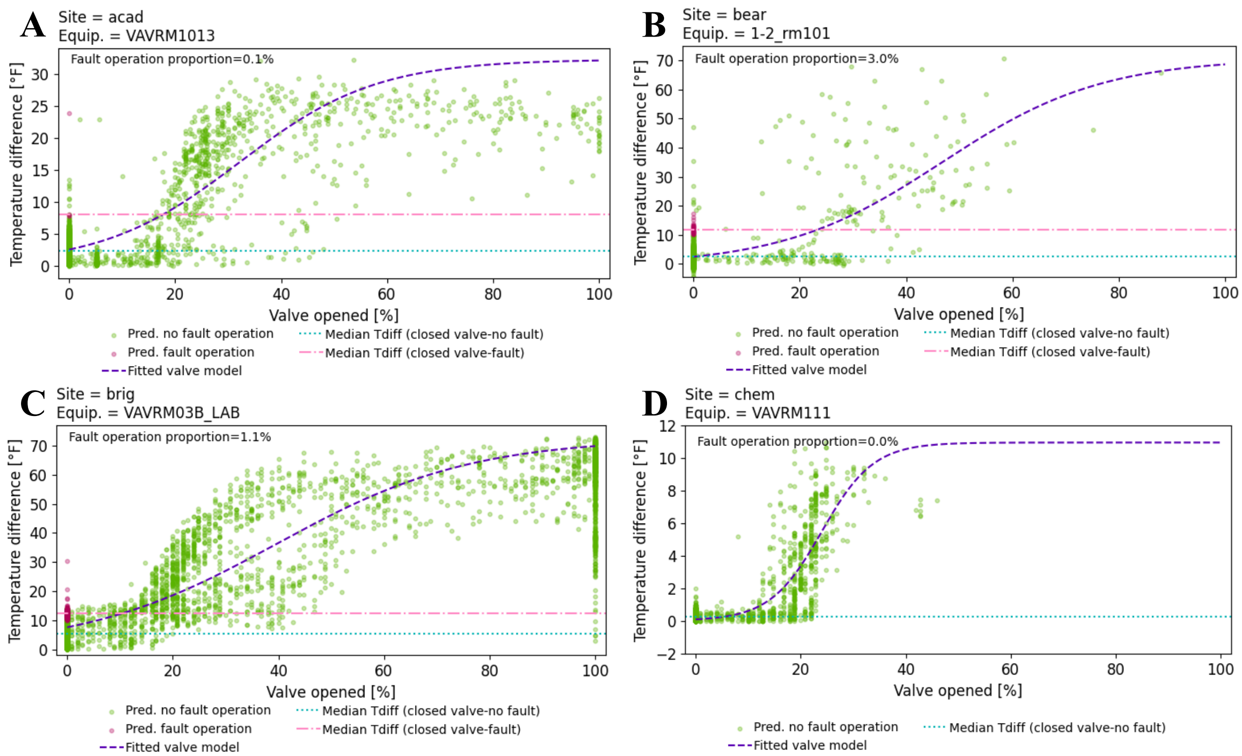

Figure 4: *Operational* data examples from a VAV unit with no fault detected and A), B) air flow rate used to filter operation mode and C), D) typical building hours used to filter operation mode.

Discussion

Correctly detecting small leakage rates in valves is typically more difficult than detecting medium and large rates (Ngo and Dexter 1999). In addition, the detection accuracy is affected when building management sensors have large uncertainty in their measurements, have drifted out of calibration, or the data collection system fails intermittently causing gaps in the datasets. We observed these issues in data as we performed the analysis. There were gaps in the analyzed datasets, we identified sensor faults that reported constant values, and we are unable to determine the uncertainty that exists in the measurements. We also do not have ground truth dataset that help with the application development. For these reasons, we chose to use conservative parameters in the application to reduce the probability of returning false positives. This includes the "accuracy parameter" for the air flow sensor measurement to determine when the VAV is in operation and assigning the passing valve temperature difference threshold in between the transition of the plateau and the linear portion of the sigmoid function model i.e., at the knee of the first bend. The temperature difference resulting from detecting a passing valve with our application would be the equivalent of opening the valve to a large percentage, around 20% as shown in Figure 4. Thus, the focus of using this application with the first iteration of parameters is to identify the most problematic VAV reheat terminal units.

We envision this application to start as a minimum set of required sensors to do a minimum set of analysis. In this case, only the supply air temperature sensors from the AHU and VAV and the valve position are required so it can be applied to many buildings. However, we can easily extend the programmatic queries to include optional sensors and Brick easily enables these types of extensions. This is the path we chose for the air flow sensor. We were able to run seven additional buildings and 802 VAV units through the application by not requiring the air flow sensor. VAV reheat systems typically have the AHU supply air temperature sensor, reheat valve position, and air flow sensors but there can be many reasons for the missing air flow sensors. For example, they are not included in the data model or data streams were not collected for various reasons. Though relatively rare, it is possible that some older buildings have pressure dependent VAV units that do not have an airflow sensor. The likely constraint on the application is the VAV supply air temperature sensor.

We can further reduce the false positives by incorporating more information into the analysis. Just like we incorporated the optional air flow sensor, we can extend the application to include more optional sensors. For example, we can add data points from the hot water plant such as boiler and pump status and hot water supply temperature to improve the usable data for our analysis. That is, we can confirm that there is hot water flowing through the pipes by knowing if the boiler and pumps are turned on. As with the relationship between AHU and VAV units, this relationship (which hot water system serves which VAV units) could be extracted automatically from the Brick model of the building. We can also add in water flow rate measurements if the sensor is available. This will enable us to get more specific in calculating an expected temperature difference as the air passes through the water coil. We can also extend the possible faults that the application can detect such as failed closed actuators or blocked/fouled coils or valves. Brick enables an automatic selection of the level of analysis as sensors are or become available in the building and the completeness of its Brick data model in including the additional information.

The application reports the magnitude in temperature rise for the passing valves and the average heat rate loss due to the fault. These metrics can be used to assess potential occupant thermal comfort issues, energy costs, and financial costs. Other metrics that faults in general need to be evaluated against include safety, environment impact, difficulty and cost of repair, and degradation rate (Katipamula and Brambley 2005). These evaluation metrics will help decision makers decide on the prioritization to correct the issue. Moreover, assessments need to be made on both the current situation of the fault and the potential effects of further degradation of the fault.

Finally, we can use the same code to extend the application to analyze any valve that controls the flow of hot or chilled water in coils that transfer heat into an air stream. All that is required is to edit the programmatic Brick query to return the required data streams, which will need to be modified to identify all water coil valves and their corresponding upstream and downstream air temperature sensors. The application evaluated in this paper focused solely on VAV reheat terminal units because there were not a significant number of AHUs or other types of equipment in the datasets with all the required data streams. Moreover, we do not know how complete the Brick models are, and we are not familiar with all the buildings in the datasets. Thus, we do not know if something else is between the upstream and downstream air temperature sensors that might affect the analysis presented in this research.

Conclusion and Future Work

In this research we developed an open-source application to detect passing valves in variable air volume (VAV) reheat terminal units. We used a framework based on the Brick schema and Mortar platform to ensure that the application is portable and can scale to many buildings. In addition, the framework allows this application to be extended to increase the performance and functionality of the application. In this first iteration, the application detects sensor faults, reversed valve operation, and passing valve faults. We analyzed 1,335 VAV reheat terminal units in 20 buildings. The application categorized 5% of VAV units with a sensor fault, 14% with a passing valve fault, and 81% with no faults. This data-driven approach makes use of the underutilized building management already present in many commercial buildings.

We applied this application to historical data, but future work could focus on the applicability and performance of the detection application when implemented in real-time. The application can be part of a suite of predictive maintenance or fault detection and diagnostic tools to help facility managers keep their building's performance at a high level. It can also be beneficial to apply the application during the commissioning process to quickly assess problems with VAV reheat terminal units that include sensor faults, reversed valve operation, and communication failures to the valve actuators. This application can easily be ported over and scaled to many buildings given that we developed it by taking advantage of the Brick schema and testing it using the variety in datasets found inside the Mortar database.

Acknowledgment

This work was supported by the US Department of Energy under grant DE-EE0008681, the California Energy Commission's Public Interest Energy Research program (contract number PIR-19-013), and the Center for the Built Environment at University of California Berkeley. We like to thank Gabe Fierro for resolving issues that enabled us to download data from Mortar platform as he is in the process of finding it a new home. Likewise, we like to thank the anonymous donors that provided data from their buildings to develop and test our application.

References

Ahmed, A., J. Ploennigs, K. Menzel, and B. Cahill. 2010. "Multi-Dimensional Building Performance Data Management for Continuous Commissioning." *Advanced Engineering Informatics* 24 (4): 466–75. Accessed July 27, 2011. https://doi.org/10.1016/j.aei.2010.06.007.

Balaji, B., A. Bhattacharya, G. Fierro, J. Gao, J. Gluck, D. Hong, A. Johansen, et al. 2016. "Brick: Towards a Unified Metadata Schema For Buildings." In *Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments*, 41–50. Palo Alto CA USA: ACM. Accessed August 18, 2020. https://doi.org/10.1145/2993422.2993577.

Balaras, C. A., and A. A. Argiriou. 2002. "Infrared Thermography for Building Diagnostics." *Energy and Buildings*, TOBUS - a European method and software for office building refurbishment, 34 (2): 171–83. Accessed January 24, 2022. https://doi.org/10.1016/S0378-7788(01)00105-0.

Bennani, I. L., A. K. Prakash, M. Zafiris, L. Paul, C. D. Duarte Roa, P. Raftery, M. Pritoni, and G. Fierro. 2021. "Query Relaxation for Portable Brick-Based Applications," 10.

EIA. 2022. "Commercial Buildings Energy Consumption Survey (CBECS) Data." Energy Information Administration, U.S. Department of Energy. Accessed January 21, 2022. https://www.eia.gov/consumption/commercial/data/2018/index.php?view=characteristics.

Fierro, G. 2021. "Self-Adapting Software for Cyberphysical Systems." Doctoral Dissertation, Berkeley, CA: University of California, Berkeley. https://www2.eecs.berkeley.edu/Pubs/TechRpts/2021/EECS-2021-159.pdf.

Fierro, G., M. Pritoni, M. Abdelbaky, D. Lengyel, J. Leyden, A. Prakash, P. Gupta, et al. 2020. “Mortar: An Open Testbed for Portable Building Analytics.” *ACM Transactions on Sensor Networks* 16 (1): 1–31. Accessed August 19, 2020. https://doi.org/10.1145/3366375.

Kaewwaewnoi, W., A. Prateepasen, and P. Kaewtrakulpong. 2010. “Investigation of the Relationship between Internal Fluid Leakage through a Valve and the Acoustic Emission Generated from the Leakage.” *Measurement* 43 (2): 274–82. Accessed January 24, 2022. https://doi.org/10.1016/j.measurement.2009.10.005.

Katipamula, S., and M. R. Brambley. 2005. “Methods for Fault Detection, Diagnostics, and Prognostics for Building Systems-A Review, Part I.” *HVAC&R Research* 11 (1): 3–25. Accessed May 14, 2021. https://www.proquest.com/docview/213538381/citation/F735847C9F784CD9PQ/1.

Lin, G., H. Kramer, and J. Granderson. 2020. “Building Fault Detection and Diagnostics: Achieved Savings, and Methods to Evaluate Algorithm Performance.” *Building and Environment* 168 (January): 106505. Accessed January 8, 2021. https://doi.org/10.1016/j.buildenv.2019.106505.

Liu, R., J. Wen, and M. S. Waring. 2014. “Improving Airflow Measurement Accuracy in VAV Terminal Units Using Flow Conditioners.” *Building and Environment* 71 (January): 81–94. Accessed March 8, 2022. https://doi.org/10.1016/j.buildenv.2013.09.015.

Miller, C. 2019. “More Buildings Make More Generalizable Models—Benchmarking Prediction Methods on Open Electrical Meter Data.” *Machine Learning and Knowledge Extraction* 1 (3): 974–93. Accessed October 27, 2021. https://doi.org/10.3390/make1030056.

Miller, C., Z. Nagy, and A. Schlueter. 2018. “A Review of Unsupervised Statistical Learning and Visual Analytics Techniques Applied to Performance Analysis of Non-Residential Buildings.” *Renewable and Sustainable Energy Reviews* 81 (January): 1365–77. Accessed March 18, 2021. https://doi.org/10.1016/j.rser.2017.05.124.

Najafi, M. 2010. “Fault Detection and Diagnosis in Building HVAC Systems.” Doctoral Dissertation, Berkeley, CA: University of California, Berkeley. https://digitalassets.lib.berkeley.edu/etd/ucb/text/Najafi_berkeley_0028E_10878.pdf.

Ngo, D., and A. L. Dexter. 1999. “A Robust Model-Based Approach to Diagnosing Faults in Air-Handling Units.” *ASHRAE Transactions* 105: 1078. Accessed May 14, 2021. https://www.proquest.com/docview/192523807/citation/ECCE1A6C6DDC47C1PQ/1.

Raftery, P., A. Geronazzo, H. Cheng, and G. Paliaga. 2018. “Quantifying Energy Losses in Hot Water Reheat Systems.” *Energy and Buildings* 179 (November): 183–99. Accessed November 14, 2018. https://doi.org/10.1016/j.enbuild.2018.09.020.

Salsbury, T. I., and R. C. Diamond. 2001. “Fault Detection in HVAC Systems Using Model-Based Feedforward Control.” *Energy and Buildings*, Special Issue: BUILDING SIMULATION’99, 33 (4): 403–15. Accessed May 13, 2021. https://doi.org/10.1016/S0378-7788(00)00122-5.

Virtanen, P., R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, et al. 2020. “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python.” *Nature Methods* 17 (3): 261–72. Accessed March 8, 2022. https://doi.org/10.1038/s41592-019-0686-2.

Zhao, X., Q. Zhang, X. Xu, Z. Shen, and B. Zhang. 2021. “A Novel Method Using Infrared Thermography for Hot Fluid Leakage Detection on Surfaces with Uneven Emissivities.”

Insight - Non-Destructive Testing and Condition Monitoring 63 (5): 273–79. Accessed January 24, 2022. https://doi.org/10.1784/insi.2021.63.5.273.