



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

What We Learned From Analyzing 18 Million Rows of Commercial Buildings' HVAC Fault Data

Eliot Crowe, Yimin Chen, Jessica Granderson¹
Hayden Reeve, Lucas Troup²
David Yuill, Yuxuan Chen³

1 Lawrence Berkeley National Laboratory
2 Pacific Northwest National Laboratory
3 University of Nebraska - Lincoln

Energy Technologies Area
August, 2022



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

Disclaimer:

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor the Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or the Regents of the University of California.

What We Learned From Analyzing 18 Million Rows of Commercial Buildings' HVAC Fault Data

*Eliot Crowe, Yimin Chen, Jessica Granderson, Lawrence Berkeley National Laboratory
Hayden Reeve, Lucas Troup, Pacific Northwest National Laboratory
David Yuill, Yuxuan Chen, University of Nebraska - Lincoln*

ABSTRACT

To achieve ambitious decarbonization goals it is critical that buildings operate to their full potential. Commercial HVAC systems, however, experience a wide range of operational faults, adversely affecting energy consumption, occupant comfort, and maintenance costs. Analytical tools such as fault detection & diagnostics (FDD) software identify and help diagnose these types of sensing, mechanical, or control-related faults.

While significant energy savings has been documented for FDD, along with limited-scale studies on technical capabilities, there is a lack of empirical data on faults being reported by FDD tools. With FDD deployment accelerating significantly over the past decade there is an opportunity to gather and analyze data on commercial HVAC operational problems at an unprecedented scale. Such data could address many questions such as: [a] What faults are most commonly reported?; and [b] How does fault reporting vary by time of year and other possible drivers?

A recent study into FDD fault reporting amassed the largest U.S. dataset of commercial HVAC air-side fault records, drawn from multi-year monitoring across over 60,000 pieces of HVAC equipment. The results of this study provide granular data on fault reporting for over 90 unique fault types. In this paper we provide an overview of the research process and highlight key findings and lessons learned. This study presents an extraordinary level of detail on FDD fault reporting characteristics across many climate zones and building types. Armed with these new insights, commercial building industry stakeholders can make better informed decisions when designing, configuring, and operating commercial HVAC systems.

Introduction

It is well-established that commercial building HVAC systems experience an array of operational and installation faults. For example, temperature sensors may drift out of calibration, damper actuators may become broken or corroded, or control sequences may not meet performance requirements. Such faults may adversely impact energy consumption, occupant comfort, or equipment life.

Past studies have documented typical energy waste from HVAC system operational faults. For example, Crowe et al. documented median whole building energy savings of 6% for implementing existing building commissioning (a quality assurance process addressing HVAC operational faults), based on a study dataset of over 1,500 buildings [Crowe et al. 2020]. Another study documented a median 9% whole building energy savings through the use of Fault Detection & Diagnostics (FDD) analytical software [Kramer et al. 2020].

Commercial buildings consume 18% of U.S. energy consumption and 36% of electricity consumption [EIA 2022]. Minimizing buildings' energy consumption continues to be a priority

for reducing carbon emissions. An additional priority is now emerging, with the emergence of the grid-enabled efficient buildings (GEB) concept [Satchwell et al. 2021]. GEB buildings are not only energy efficient, but they also need to be reliable grid assets with the ability to dispatch load when required, to allow for better management and reliability of generation resources. Under this paradigm the impact of leaking water valves, faulty sensors, broken dampers, and control sequence conflicts may be magnified if it results in a building falling short of grid requirements at a time of critical need. For example, a leaking chilled water valve would limit the ability to temporarily adjust zone temperatures upward in response to a demand response signal, thereby reducing load reduction potential.

Understanding the energy savings potential from resolving HVAC operational issues has helped to build the business case for commissioning and analytical software (energy management and information systems, or EMIS). While energy savings are now well documented, the nature and frequency of operational faults discovered in commercial buildings has been less studied. Kim et al. carried out a comprehensive review on HVAC fault prevalence [Kim et al. 2021]. In the research, a literature review of 26 studies, along with 25 expert interviews were employed to fully understand the current state of knowledge, gaps and potential value of research on HVAC fault prevalence in commercial buildings. Through the study, the authors uncovered unmet needs, gaps and corresponding suggestions. For example, the authors found little research to fully document the prevalence of HVAC faults at the desired level of granularity, consistency, and scale. In addition, there is a need to resolve HVAC fault data under various drivers such as building type, system type and climate zones, so that fault impacts can be quantified and FDD approaches can be efficiently developed. Based on the findings, the authors suggested a complete study on commercial buildings' HVAC fault prevalence to bridge knowledge gaps.

This paper reports a summary subset of results for a 3-year study funded by the United States Department of Energy (U.S. DOE). The purpose of the study was to gain a more granular understanding of the relative levels of prevalence for a wide range of commercial HVAC fault types. Obtaining granular data on building performance across a large population is notoriously challenging and expensive, hence the study team focused on obtaining HVAC fault data as reported by FDD tools. This enabled the collection of a large fault dataset from a relatively small number of sources (i.e., several FDD software vendors could each provide over a year of fault data across tens or hundreds of buildings).

In this paper we describe the data collection method, normalization of data to common time resolution and fault naming conventions, fault prevalence metrics, and a selection of key results for air handling units (AHUs) and air terminal units (ATUs); packaged rooftop HVAC unit (RTU) fault data analysis is ongoing, and excluded from this paper. We then discuss the potential sources of uncertainty in the results, and implications for HVAC operators, system designers, technology developers, and researchers. We conclude with final recommendations.

Research Method

The HVAC fault prevalence study involved multiple steps to obtain fault data, clean and normalize that data, develop reporting metrics, and analyze results (Figure 1). Each of these steps is described in detail below.




Figure 1. Fault Prevalence Study Steps

Obtaining and Cleaning Study Data

Study data were drawn from six FDD software vendors and one large building portfolio owner. The scope of the data request was limited to built-up air handlers (AHUs) and associated air terminal units (ATUs), and packaged rooftop HVAC units (RTUs). In total, we received data covering 3,660 AHUs, 53,865 ATUs, and 7,974 RTUs. In almost all cases we received at least twelve months' worth of reported faults, and in some cases over 24 months. In addition to receiving data on reported faults we also received building metadata summarizing (at minimum) building type and building location. Metadata included full AHU/ATU equipment listings including equipment for which there were no reported faults. Over the course of several weeks or months we worked with data providers to gain a clear understanding of the data shared, for example:

- Clarifying fault naming definitions;
- Understanding linkages between different buildings and the inter-relationships between equipment;
- Addressing data gaps and inconsistencies;
- Clarifying date ranges over which buildings were monitored (i.e., the first and last dates of reported faults in a building are not necessarily the full monitoring time period);
- Clarifying, where possible, the presence or absence of non-standard components on specific buildings (e.g., a return air carbon dioxide sensor may not be present on all AHUs).

Once the data were fully reviewed for accuracy, fault names and time intervals were standardized.

Applying HVAC Fault Taxonomy and Standard Time Resolution

Applying a consistent fault naming structure is the key to analyzing data from disparate sources. The data providers for the fault prevalence study all used unique fault naming approaches, and in some cases their datasets included several fault names for the same fault (e.g., they allowed FDD users to rename faults when initially configuring the software). For datasets covering AHUs/ATUs there were many hundreds of unique fault names, up to 1,105 in one case. Through manual review, all raw fault names (i.e., names provided in datasets received from partners) were checked and, where applicable, mapped to standardized fault names based on an HVAC fault taxonomy.

The HVAC fault taxonomy (described in [Chen et al. 2020]) classifies HVAC faults on four levels, each represented by a different section in a unique fault identification (ID) code:

- Equipment type (currently limited to AHUs, ATUs, and RTUs)
- Fault location in the system (e.g., supply air, mixed air, etc.)
- Component or parameter (an example component is a temperature sensor, an example parameter is measured static pressure)
- Fault mode (e.g., frozen, stuck, abnormal)

An example fault ID takes the form “AHU-Supply_air-Temperature_sensor-Frozen.” The taxonomy design allows for both condition-based (CB) faults and behavior-based (BB) faults. CB fault reporting specifies the component where the fault symptom is observed (e.g., return air temperature sensor frozen), whereas BB faults report the observed behavior without pinpointing the specific component (for example, ‘supply air temperature abnormal’ behavior could be due to a faulty sensor reading or the air temperature could truly be too high or low). Further detail on CB and BB fault definitions is provided in Kim et al. 2021. Using this fault ID structure it is possible to analyze data at various levels of aggregation. For example, the fault ID “AHU-Supply_air-Temperature_sensor,” would be applied for reporting the incidence of any of the possible faults on that specific component, regardless of the difference fault modes that may have occurred (e.g., sensor frozen, biased, etc.)

Through manual review, and communication with data providers where needed for clarification, all in-scope faults were mapped to fault names in the HVAC fault taxonomy. In some cases a fault’s mapped ID code could not be fully resolved; for example, in a case where a faulty component was reported but not its fault mode the ID includes “NA” (RTU-Supply_air-Temperature_sensor-NA). From an initial total of 2,798 raw fault names, we ended up with 275 mapped fault names for the study (in many cases multiple raw fault names were translated to a single mapped fault name for the study). Raw fault names that were not mapped included out-of-scope equipment such as boilers, and out-of-scope faults such as abnormal fan current. Figure 2 illustrates the full scope of system locations, components/parameters, and fault modes represented in the study dataset.

As noted above, the taxonomy provides a unique fault ID with four separate sections. This allows for analysis at different levels of aggregation. Analyzing unique faults separately results in a long list, but results may be analyzed and reported at the level of component or system location. For example, for simplicity one may want to initially aggregate analysis to the level of component, reporting the ‘faultiness’ of different components before drilling down to the extra detail of specific fault modes.

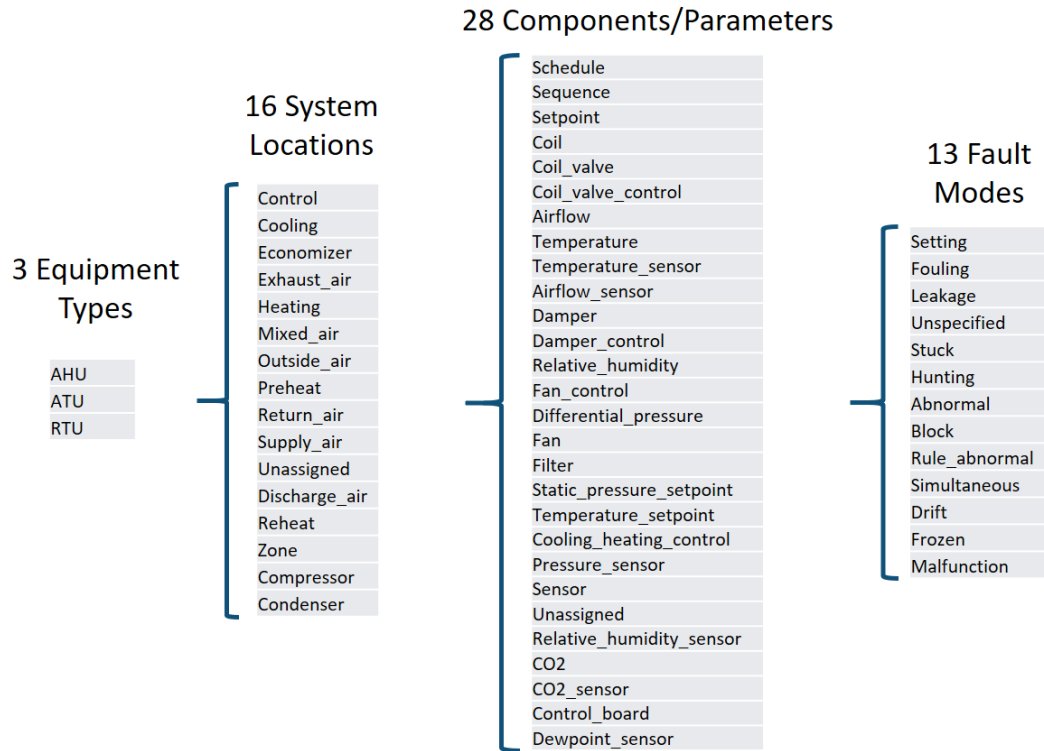


Figure 2. Fault naming parameters included in the study, based on HVAC fault taxonomy structure

Concurrently with converting raw fault data to the taxonomy naming, fault reporting was converted to a common daily resolution. Raw fault data was provided with different time resolutions, in some cases documenting dates on which a fault was reported, sometimes noting a start date and duration, and in either case time of day might also be reported. For the study, all data were converted to ‘binary daily format’ (BDF), whereby a row in the fault database represented a unique fault occurring on a specific piece of equipment on a specific date, irrespective of whether the fault occurred for 1 hour or 24 hours on that specific day. Each row of the database, therefore, may be defined as a “Fault_Day,” and if multiple different faults occur on the same piece of equipment on a specific date it would represent multiple Fault_Days. Using this approach, the total number of Fault_Days collected in the study database (i.e., the number of rows of data) was almost 18 million.

Data Analysis and Metrics

Based on the size and granularity of the study dataset there are many metrics and analysis approaches that may be developed. To date five metrics have been developed; two are described and reported here, both derived from the same ‘Percent Time Faulted’ (PTF) analysis. Figure 3 describes the calculation approach to determine PTF. For a specific fault type, a percentage value is calculated for every piece of equipment, representing the percent of its monitored days¹ where

¹ Building metadata defines the start and end date of the FDD monitoring period applicable to all equipment in that building, from which the number of monitored days is calculated; dates on which a specific fault occurs are derived from the binary daily fault (BDF) database.

the fault was reported. In almost all cases equipment includes 365 or more monitored days, though Figure 3 shows a 5-day time-series to illustrate the calculation approach.




Figure 3. Percent Time Faulted (PTF) calculation concept (“1” represents a day on which a fault was reported)

Using the calculation approach shown in Figure 3, a PTF value can be calculated for every fault for every piece of equipment, e.g., a single piece of equipment will have a PTF value for every fault type the FDD software was programmed to report in that building. It is then possible to establish a Mean Percent Time Faulted (MPTF) value for every fault type, allowing for the ranking of fault types from most to least common. Preliminary MPTF analysis uncovered a challenge, in that the MPTF value was affected by two different factors:

1. The portion of the population of equipment that experienced a given fault; and
2. For the equipment that experienced a given fault, the portion of their time series saw the fault’s presence

To illustrate these two factors we can consider two extreme cases: [1] A fault occurs on all AHUs, and is present for 50% of the time; and [2] A fault occurs on half of the equipment, and for those pieces of equipment it is present for 100% of the time. In both cases the MPTF would be the same (50%), though the fault prevalence characteristics are very different. For this reason, MPTF was superseded by two related metrics, described in Table 1.

Table 1. Metric descriptions

Metric	Description
Pct_Affected	Percent of the population of a given equipment type that experienced a given fault on at least one day in its monitored time period. A measure of how commonly a fault is reported
MPTF_Affected	MPTF for those equipment in the Pct_Affected data subset (i.e., excluding all equipment that never saw a given fault). A measure of how frequently one might expect a fault to be reported on a single piece of equipment.

In addition to reporting the metrics listed in Table 1, the results below also summarize the volume of faults reported each month per building, per AHU, and per ATU (for faults within scope for this study). At time of writing, RTU faults analysis is still being refined, and is not included here.

Results

Table 2 shows the average number of faults reported by FDD tools per month (denotes unique faults reported; each of those faults may be reported multiple times within a given month). On a per-building basis the range is very wide, since buildings vary considerably in the amount of installed equipment. For example, the number of pieces of monitored equipment in a building ranged from 1 to 1,778. While there is significant variation in per building results, an interquartile range of 27 to 274 provides an illustration of the volume of reported faults that a building's operational team would need to consider (especially as many operational teams oversee multiple buildings across a campus or portfolio, and these study results only cover AHUs and ATUs). These results confirm the risk of potential data overload for operations staff, and emphasizes the need for tools and processes to manage and prioritize responses to reported faults.

When data are divided into AHU and ATU faults it is possible to normalize per unit of equipment, allowing for comparison of fault reporting prevalence. Table 2 illustrates the difference in typical number of faults for AHUs and ATUs; on average, each AHU has 3 different faults reported each month (often being reported across multiple days in a given month), whereas ATUs average just over 1 fault per month. AHUs are more complex equipment and tend to have more applicable fault types (26 on average, based on faults mapped to the study taxonomy), compared to ATUs (average 13 mapped faults). While ATUs have fewer faults per unit (average 1.2 compared to 3.0), the total number of reported faults is higher for ATUs, as the study dataset includes an average 15 ATUs for each AHU. Across the whole study dataset there are 11 million Fault_Days for ATUs and 2 million Fault_Days for AHUs.

Table 2. Average number of reported faults per month

	Average Reported Faults per Month	Sample size
Per Building	245	317 buildings
Per AHU	3.0	3,660 AHUs
Per ATU	1.2	53,865 ATUs

The volume of reported faults can vary significantly between AHUs and between ATUs, with some equipment reporting zero faults (within the scope of faults included in this study). Looking at the number of Fault_Days reported for every individual piece of equipment across the whole study dataset, we find that the top 34% of AHUs (i.e., the top 34% of AHUs when sorted by highest number of Fault_Days) account for 80% of total AHU Fault_Days in the study dataset. By coincidence, it is also the case that the top 34% of ATUs account for 80% of total ATU Fault_Days. The relationship between volume of faults for AHUs and ATUs can vary within a given building, however, with Figures 4 and 5 providing an example. Figure 4 shows a relatively similar number of Fault_Days for the six AHUs in this example building. Figure 5 shows the Fault_Days for each ATU in the same building as used in Figure 4. Here we see a very wide distribution in the volume of faults reported, with 20 ATUs accounting for 80% of all Fault_Days, and 16 ATUs reporting no faults.




Figure 4. Number of Fault_Days for all AHUs in Building 17




Figure 5. Number of Fault_Days for all ATUs in Building 17

Moving from quantity of faults reported to looking at individual faults, Figure 6 shows the top reported fault types for AHUs, based on the percent of the population of AHUs experiencing a given fault at least once (“Pct_Affected”). These faults are aggregated to one level, i.e., they are defined either by the specific component on which they occur (e.g., supply air temperature sensor) for CB faults or by a measured parameter (e.g., zone temperature) for BB fault types, and fault mode (e.g., frozen, abnormal, etc.) is not specified. Figure 6 shows the 21 fault types occurring on at least 20% of AHUs (21 fault types are shown, out of a total of 69 AHU fault types).




Figure 6. Most commonly reported faults for AHUs, based on Pct_Affected value (chart shows faults 20% or higher)

Observations from Figure 6 include:

- The most highly reported fault relates to supply air temperature (SAT) setpoint, appearing on 55% of AHUs at some point, being the only fault type to occur on over half the AHUs in the study dataset.
- While the three most common faults are behavior-based (BB), only one other fault shown in Figure 6 is BB (“cooling-cooling,” where cooling capacity is not meeting requirements). The remaining 17 faults are condition-based (CB) faults.
- The components most commonly affected by faults are sensors (5), valves (3), and dampers (3).
- The most common locations for faults to be reported are the supply air section (5) or related to the cooling coil (3).

Figure 7 illustrates the most commonly reported faults for ATUs. As above, Figure 7 shows the faults occurring on at least 20% of ATUs (8 fault types are shown, out of a total of 21 ATU fault types).




Figure 7. Most commonly reported faults for ATUs, based on Pct_Affected value. (chart shows faults 20% or higher)

Observations from Figure 7 include:

- The most highly reported faults relate to faulty zone temperature sensors or abnormal zone temperature, appearing on 46% of ATUs at some point
- Aside from zone temperature-related faults, no other fault occurs on over 34% of ATUs
- The eight most commonly reported faults are evenly split between CB and BB faults

Figures 6 and 7 indicate the percent of equipment for which certain faults are reported, a measure of how common a fault is. In contrast, Figures 8 and 9 show the Mean Percent Time Faulted (MPTF_Affected) for AHUs and ATUs respectively, indicating the average *portion of time* a fault is reported (e.g., an AHU may have a return air temperature fault reported on 23% of the days across its entire time series). MPTF_Affected is a measure of how persistently/frequently a fault is reported over time, and is calculated only for those pieces of equipment on which a given fault occurred at some point. Figure 8 shows the 18 AHU faults with an MPTF_Affected value of 20% or greater.




Figure 8. Mean Percent Time Faulted, for AHUs experiencing a given fault at least once (MPTF_Affected). Excludes AHUs that never saw a given fault. Chart shows faults with MPTF of at least 20%

Observations from Figure 8 include:

- Supply air relative humidity faults are reported over the longest periods, close to half of the time (43%) on average
- Three of the top six faults shown in Figure 8 occur in AHUs' return air section, two are related to supply air
- Nine faults are reported more than a quarter of the time that equipment is monitored.
- The 18 faults being reported at least 20% of the time are evenly split between CB and BB fault types (in contrast to Pct_Affected, where the majority of the top faults were CB)

Six AHU fault types occur on at least 20% of equipment *and* have those faults occur at least 20% of the time (i.e., they are shown on Figures 6 and 8): supply air temperature setpoint, supply air static pressure setpoint, economizer sequence, supply air temperature sensor, coil valve², and airflow sensor³.

Figure 9 shows the 13 AHU-related faults for which MPTF_Affected is 20% or greater.

² Based on the data provided the exact type/location of coil valve and airflow sensor could not be determined




Figure 9. Mean Percent Time Faulted, for ATUs experiencing a given fault at least once (MPTF_Affected). Excludes ATUs that never saw a given fault

Observations from Figure 9 include:

- Three ATU faults are reported more than 50% of the time: Zone CO₂ sensor, Zone RH, and cooling related faults. All other faults in Figure 9 are reported 21% - 29% of the time
- Three of the top four faults shown in Figure 9 are BB faults, but BB faults make up the minority of faults overall (5 of the 13 shown)
- Figure 9 shows results for 13 fault types, the majority of the total 20 ATU fault types covered in the study dataset

Five ATU fault types occur on at least 20% of equipment *and* have those faults occur at least 20% of the time (i.e., they are shown on Figures 7 and 9): Zone temperature, reheat coil valve, discharge air temperature sensor, heating capacity faults, and faults relating to control sequences.

The metrics developed in this study allow for analysis of fault likelihood by month of year, which can be useful in understanding the impact of weather as a fault driver. Monthly variation can be assessed using the coefficient of variation (COV), which is the standard deviation of the percent fault presence values across the twelve months of the year divided by the mean value across those twelve months. Sorting individual fault types by coefficient of variation allows for a ranking of faults based on the degree of temporal variation (where a higher COV denotes that fault presence varies more significantly by month). Figure 10 shows the COV for all AHU faults for a single data partner (left), and also illustrates monthly presence values for two individual fault types, one with high COV (top right) and one with low COV (bottom right).




Figure 10. Coefficient of variation (COV) for all AHU faults reported by a single data partner (left), and charts showing monthly fault presence values for two example fault types, one with high COV (top right) and one with low COV (bottom right)

COV analysis based on month of year is ongoing, and similar analysis is planned for assessing the variation in fault reporting by building type and by climate zone.

Discussion

Study results confirmed anecdotal evidence that taking a proactive approach to uncovering HVAC operational issues can result in a very high volume of reported faults. An average of 245 reported faults per month per building (in many cases being reported on multiple days across that month) is a significant workload to review, interpret, and prioritize. The data overload challenge is even more pronounced when considering that [a] There is other HVAC equipment beyond AHUs and ATUs to be monitored; and [b] It is very common for campuses/portfolios to consolidate FDD reporting across many buildings to a central energy management team. The high volume of reported faults emphasizes several key needs:

1. Methods to prioritize faults by impact type (energy, comfort, and/or maintenance) and impact magnitude. This is available on some tools and is the subject of some ongoing

studies [Chen et al. 2021], and is worthy of further study, development, and standardization.

2. Improved root cause diagnostics. Many reported faults require additional manual investigation to confirm they are true faults and understand root cause (particularly behavior-based faults). Efforts to improve root cause diagnostics, even for a subset of the most common faults, could significantly improve the turnaround time for fault resolution.
3. Automating the correction of faults. This is an area of current study [Pritoni et al. 2022]; as above, even if this could apply to only a subset of faults it would improve buildings' efficiency and reduce the burden on operations staff.

This study gathered data at a scale and granularity beyond similar prior studies found in literature, but like all similar studies there are several sources of uncertainty in results:

- Some reported faults may be false positives. For example, a data communication issue may be reported as a frozen sensor reading.
- A single root cause may be represented in several faults. For example, an outdoor air damper stuck open could cause faults that are reported as an economizer fault and a mixed air temperature fault.
- There may be false negatives in the study dataset. For example, AHUs never reporting a frozen CO₂ sensor may not have a CO₂ sensor installed. In other cases, a fault may exist but its symptoms may not manifest across all days in the study dataset, e.g., an economizer damper stuck closed may not be reported on hot days when being closed is the correct position.
- Differences in fault detection algorithms between data partners could result in a different volume of fault records being triggered for similar operating conditions between buildings.
- It is possible that discrepancies in how raw fault names were interpreted and translated into taxonomy naming conventions affects the relative prevalence between faults.

The study team collaborated with data partners over an extended period to address these sources of uncertainty to the greatest extent possible, and developed quality control analysis tools and visualizations to catch inconsistencies where possible. We are also in progress with more detailed onsite assessments of true positives and false negatives for a small sample of the study dataset buildings, and tracking related ongoing research to objectively compare detection sensitivity between different FDD tools.

It should also be noted that even though the study dataset is large and covers many building types and climate zones, it is a limited sample and not demonstrated to be representative of the U.S. commercial building stock. It may also be subject to availability bias; for example, building owners who have chosen to invest in FDD deployment may be more likely to own buildings that perform above average. If true this would mean that commercial buildings in general would expect an even higher volume of reported faults than that shown by this study.

Despite these sources of uncertainty the study provides a level of granularity that helps in understanding how reported faults break down in terms of condition-based vs. behavior-based characteristics, the system locations or components in which they most commonly occur, and comparing prevalence between sensor faults, control-related faults, and mechanical faults.

While the scope of this study did not attempt to quantify the energy savings impact of faults, the volume and nature of the faults reported in the study reinforce the credibility of prior literature on the savings benefits of FDD (9%, as reported in Kramer et al. 2021) and existing building commissioning (6%, as reported in Crowe et al. 2020).

The high incidence of sensor-related faults highlights a priority for improvement. This is especially needed given the potential increase in building system sensing as more Internet of Things (IoT) devices are installed. For sensor manufacturers, further study on the root causes of sensor faults is warranted. For system designers, reviewing design approaches for ATUs to address the ‘problem minority’ could have significant benefits for both occupant comfort in high-fault zones, and also potentially greater energy savings (since a small number of ATUs not meeting minimum airflow requirements can result in system adjustments across all ATUs, with potential to significantly drive overall energy consumption upward).

Another area for further study could look at fault persistence. This study did not look into the possible causes of faults lingering for long periods of time, e.g., whether a building operations team did not have the resources to diagnose and fix faults, or the fault report was considered a nuisance or false alarm, or if a fault was recurring intermittently/frequently. Understanding these factors, and the performance/savings implications of persistent faults, would help in determining resolutions that improve operations teams’ resource management and bring buildings into better control.

Conclusions and Future Research

Aggressive decarbonization goals are driving the need for more energy efficient, grid-responsive buildings. Advanced building controls are a key part of realizing GEB performance objectives, but [a] they are not a ‘fit and forget’ solution, and [b] buildings need to be properly under control for GEB strategies to be implemented effectively. For example, how much more effective could a global zone temperature setpoint adjustment be if zone temperature sensors were not reported faulty 27% of the time? It is critical to:

- Identify and address the design/operational root causes of HVAC system faults so that they can be avoided,
- Continue to support building analytics deployment as an essential tool for achieving projected energy benefits of GEB buildings and advanced controls in general, and
- Improve FDD analytics software to support better and more standardized fault identification, root cause diagnosis, and potential to auto-correct faults, in order to reduce resolution time for faults when they do occur

While this study was significant in scale and resolution there is potential to go further, in several dimensions. Continuing to build consistent fault data across more building types and climate zones will support greater statistical significance of results. Gathering data over a longer period of time, and with additional metadata, would support a deeper understanding of the long term persistence/recurrence of faults and contributing factors. The study approach could also be extended to other HVAC equipment such as chillers, boilers, pumps, and cooling towers. Any future work would also benefit from greater standardization of fault naming and hierarchical structure, and also moving toward more consistency in setting fault trigger thresholds.

FDD implementation has been shown to reduce building energy consumption by 9%; it is assumed that a significant portion of the buildings covered by this study have already made

significant energy improvements, and yet are still seeing a high volume of faults. This would suggest that there is still significant energy savings potential beyond the 9% if we can more aggressively address building performance issues.

Acknowledgements

This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Building Technologies Office, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. The authors are grateful to Erika Gupta and Brian Walker at the U.S. Department of Energy for their support of this project.

References

- Chen, Yimin, Eliot Crowe, Guanjing Lin, and Jessica Granderson. "What's in a Name? Developing a Standardized Taxonomy for HVAC System Faults." Proceedings of the 2020 ACEEE Summer Study on Energy Efficiency in Buildings (2020). DOI: 10.20357/B7W02S
- Chen, Yimin, Guanjing Lin, Zhelun Chen, Jin Wen, Jessica Granderson. "A simulation-based evaluation of fan coil units fault effects." Energy and Buildings 263 (2022). DOI: <https://doi.org/10.1016/j.enbuild.2022.112041>
- Crowe, Eliot, Evan Mills, Tom Poeling, Claire Curtin, Diana Bjornskov, Liz Fischer, and Jessica Granderson. "Building Commissioning Costs and Savings Across Three Decades and 1,500 North American Buildings." Energy and Buildings 227.110408 (2020). DOI: 10.1016/j.enbuild.2020.110408
- Energy Information Administration. "EIA Monthly Energy Review, March 2022." EIA (2022)
- Kim, Janghyun, Kim Trenbath, Jessica Granderson, Yimin Chen, Eliot Crowe, Hayden Reeve, Sarah Newman & Paul Ehrlich. "Research challenges and directions in HVAC fault prevalence." Science and Technology for the Built Environment (2021). DOI: 10.1080/23744731.2021.1898243
- Kramer, Hannah, Claire Curtin, Guanjing Lin, Eliot Crowe, and Jessica Granderson. "Proving the Business Case for Building Analytics." Lawrence Berkeley National Laboratory (2020). DOI: 10.20357/B7G022
- Pritoni, Marco, Guanjing Lin, Yiming Chen, Raphael Vitti, Christopher Weyandt, Jessica Granderson. "From fault-detection to automated fault correction: a field study." Building and Environment, 214 (2022).
- Satchwell, Andrew, Mary Ann Piette, Aditya Khandekar, Jessica Granderson, Natalie Mims Frick, Ryan Hledik, Ahmad Faruqui, Long Lam, Stephanie Ross, Jesse Cohen, Kitty Wang, Daniela Urigwe, Dan Delurey, Monica Neukomm, and David Nemtzow. "A National Roadmap for Grid-Interactive Efficient Buildings." United States Department of Energy (2021).