The cost of HVAC demand response: Using experimental data to track down the causes of inefficiency in sub-hourly HVAC load shifting

Austin J. Lin, University of Michigan
Johanna L. Mathieu, University of Michigan

ABSTRACT

Control of Heating, Ventilation, and Air Conditioning (HVAC) systems is a promising method for short timescale (sub-hourly) demand response (DR). The thermal mass of a building allows for flexibility in HVAC power consumption without causing significant changes in building temperatures, enabling the building to act like a battery and the grid to integrate more renewables. For effective real-world implementation, the tradeoffs between grid benefits and impacts on building services need to be understood. Previous studies have shown that short timescale DR events targeting HVAC fan load shifting can cause buildings to consume excess energy when returning to normal operation. Past studies do not agree on the magnitude of this excess energy consumption, with significant differences between experimental and simulation results. Past studies have also not fully explained the reasons for this excess energy consumption. In this paper, we leverage a combination of simulation and experimental data from buildings in which open-loop global thermostat adjustment sub-hourly load shifting was conducted. We identify several potential factors contributing to excess energy consumption. Unlike past work, which only investigated changes in room temperature, we investigate changes in ventilation. Ultimately, our experimental data are not enough to determine the extent of inefficiency from each identified source; however, our findings indicate that at least part of the excess energy consumption is a byproduct of building control design. We provide a roadmap for future experimentation to determine the extent of inefficiency and how excess energy consumption may be reduced by changing building control design.

Introduction

Commercial buildings Heating, Ventilation and Air Conditioning (HVAC) systems offer a large potential resource for providing Demand Response (DR) to power grids. DR is the action of a load on the power grid to change its consumption to help balance power generation. This becomes increasingly important as power grids shift away from traditional generation towards renewable resources, from which power generation can be intermittent and uncertain (Taylor et al. 2016). HVAC systems are well-suited for DR as building thermal mass allows for HVAC power consumption to be adjusted without resulting in immediate changes to building temperature (Hao, et al. 2014). Commercial building HVAC DR schemes can vary in timescale, with some strategies targeting full-day planning (Liu et al. 2023) and others targeting fast timescale services such as frequency regulation (Vrettos et al. 2018).

A specific type of DR is load shifting, where the goal is to temporally shift power without changing the building's overall energy consumption. This is traditionally done by shifting load from on-peak hours to off-peak hours, but here we consider load shifting within a single hour. This type of load shifting can be valuable for grid balancing, especially at higher levels of renewable penetration. HVAC systems have the ability to perform load shifting on sub-hourly timescales (Beil et al. 2015) and can be controlled through Global Thermostat Adjustment (GTA) (Keskar et al. 2022). The buildings under consideration use a Variable Air Volume (VAV) terminal reheat system where an Air Handling Unit (AHU) supplies air to a central duct. With GTA, the temperature setpoint in each room is modified slightly to trigger changes in building cooling, which leads to changes in fan (and chiller) power consumption.

It is common in the power systems literature to view some specific DR resources as virtual batteries (Raman and Barooah 2020), where the battery analogy compares the power and energy consumption during the event to the counterfactual baseline power and energy consumption if no event had occurred. Prior work has found that the AHU fans in VAV HVAC systems behave as inefficient virtual batteries, consuming more energy than otherwise would have been consumed during a settling period when the building returns to normal operation after a load shifting event (Beil et al. 2015). Later experimentation found that, in certain cases, AHU fans can also consume less energy than the baseline energy (Keskar et al. 2020). These changes in energy consumption due to load shifting make HVAC systems "non-ideal batteries," and these nonidealities, affecting the value and cost of DR, should be considered by both the building operator and the grid operator. We acknowledge that buildings have more complex physics and controls than batteries (Afroz et al. 2018), and so some researchers and practitioners have questioned the usefulness of the battery analogy. However, the analogy continues to be used by the power systems community as it provides a pathway to simplify the integration of extremely large numbers of distributed energy resources (including grid-interactive buildings) by representing these complex, high-dimensional systems with approximate low-dimensional models.

Throughout previous studies, there is no consensus on the magnitude of the excess energy consumption and whether energy is over- or under-consumed. The experimental results presented by Beil et al. (2015) showed much more excess energy consumption than results presented by Keskar et al. (2020) and Lin et al. (2024b). In frequency regulation experiments, there was no observed change in energy consumption by HVAC fans (Vrettos et al. 2018). These experimental results do not agree with simulations performed by Lin et al. (2017) that explain the excess energy consumption through changes in room temperature. Specifically, they do not agree with the conditions for over- or under-consumption of energy compared to the baseline. Raman and Barooah (2020) later showed that successive load shifting events result in less energy over- or under-consumption. Overall, there is a lack of consensus on VAV HVAC load shifting inefficiency – both its magnitude and causes. The factors that cause load shifting inefficiency in individual buildings and, importantly, the factors that cause the differences in load shifting inefficiency between buildings are not yet fully understood.

From past sub-hourly load shifting experiments and simulations it is unclear whether the excess energy consumption is inherent to the building or an unintended byproduct of the building control design or tuning. If the latter is the dominant factor, it may be possible to reduce excess energy consumption through control design. In any case, it is important for us to better understand the building physics and controls causing excess energy consumption such that we can properly value short timescale HVAC load shifting and compensate participants for this value (e.g., cover the cost of excess energy consumption when DR provides a net positive benefit to the building-grid system).

In this paper, we use an iterative data-first method to identify several potential factors contributing to excess energy consumption by AHU fans in VAV HVAC systems during sub-hourly load shifting. We examine a mix of experimental and simulation data. We investigate how building operating points change during load shifting and find, unsurprisingly, that room temperature alone does not completely explain excess energy consumption. We identify three potential factors contributing to load shifting inefficiency relating to airflow and pressure: 1) pressure reset control in the building may be triggered by events, 2) fast fluctuations in airflow triggered by events may be impacting duct pressure sensors, and 3) economizer operation may be impacted by events, changing the total amount of fresh air drawn into the building. Additionally, we discuss the limitations of our experimental dataset that may be influencing our perception of the inefficiency. We conduct building simulations leveraging the Modelica Buildings Library (Wetter et al. 2014) to show that excess energy consumption could be reduced through changes to building control design, specifically, disabling the pressure reset control. However, due to the limitations of the experimental dataset, we are unable to demonstrate this phenomenon in practice, and so the extent and the potential for reduction of the excess energy consumption remain unclear. We then provide a summary of the additional experimentation needed to determine the extent of the excess energy consumption and whether excess energy consumption can be reduced through changes to building control design.

The rest of this paper is organized as follows. We present a background section discussing past explanations for excess energy consumption, a methods section outlining our analysis process and data, and a results section providing our key findings and summary of suggested future experimentation. Finally, we conclude the paper.

Background

In this section, we provide background information about the nature of excess energy consumption of AHU fans during short timescale VAV HVAC load shifting and outline the previous explanations that rely on changes to room temperature.

There is not a clear consensus on the magnitude of and reasons for the changes to building energy consumption caused by sub-hourly HVAC load shifting (MacDonald et al. 2020). The direction of load shifting, which refers to if power is first increased or decreased relative to the baseline power, influences event energy consumption. An UP-DOWN event is one in which the power is first increased (through a reduction in temperature setpoint via GTA) and then decreased (through an increase in temperature setpoint via GTA), while a DOWN-UP event is one in which power is first decreased and then increased. Past work has used 30-minute or 1-hour events, with setpoints persisting for half of this period, i.e., a 1-hour event increases (decreases) temperature setpoints for the first 30 minutes and then decreases (increases) temperature setpoint for the second 30 minutes to achieve a load shift in which the HVAC system consumes the same amount of energy across the event as the baseline energy consumption. Past experimentation has found that setpoint changes within 2°F of the user setpoint have no effect on occupant comfort (Keskar et al. 2022). Additionally, past experimentation has found that short timescale load shifting comes primarily from changes in AHU fan power consumption, though there may also be some impact on chiller power consumption (Keskar et al. 2022). We focus our work on changes to AHU fan power consumption in response to the temperature setpoint changes via GTA.

Figure 1 shows an illustration of the changes in fan power and average room temperature during an UP-DOWN load shifting event. This illustration shows how the fan power and average room temperature might change in response to temperature setpoint changes. After the event ends, there are two possibilities for fan power returning to normal operation: overshooting the baseline power (case 1 in Figure 1) or undershooting the baseline power (cases 2 & 3 in Figure

1). If the power follows case 1, the event causes the building to consume excess energy to return to normal operation. If instead the power follows case 2, the building under-consumes energy after the event. Experimental results, such as those from Keskar et al. (2020), found that fan power tended to follow case 2. This directly contrasts the simulation results of Lin et al. (2017) where fan power followed case 1. The DOWN-UP events are the reverse of the UP-DOWN events (and can be visualized by mirroring Figure 1 along the time axis), where the experimental results tended to follow case 2 (which is now consistent with over-consumption) and the simulation results follow case 1 (which is now consistent with under-consumption). The main difference between the experimental results and simulation results is the presence or lack of overshoot, i.e., when the fan power consumption crosses the baseline fan power consumption during the post-event settling period (case 1 in Figure 1).

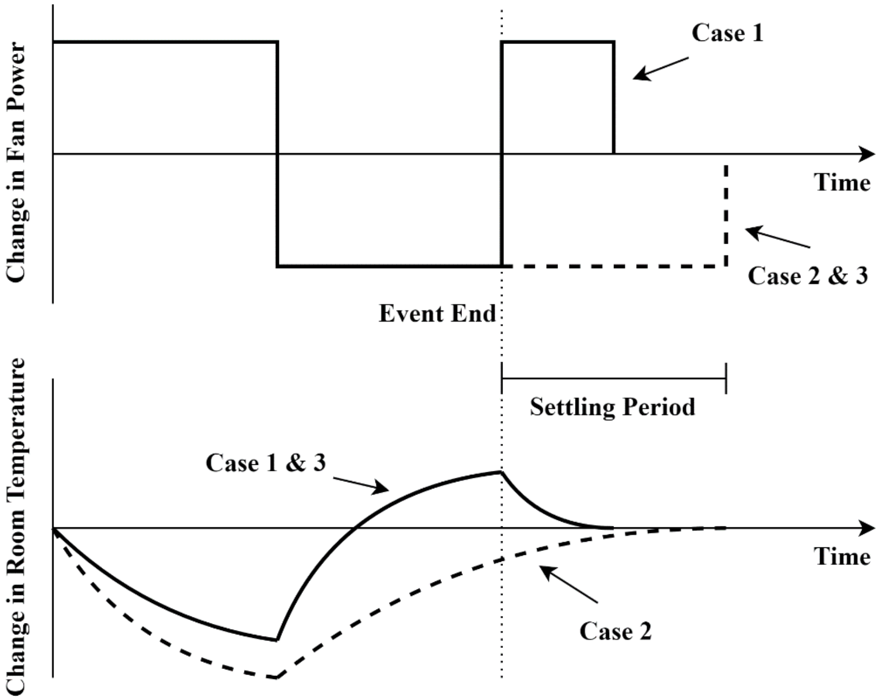

Figure 1: Illustration of how an UP-DOWN load shifting event (i.e., a sequence of temperature setpoint changes via GTA) can change AHU fan power and average room temperature. Three cases are presented. Case 1: overshoot in fan power and room temperature crosses the original temperature setpoint (i.e., the temperature baseline). Case 2: no overshoot in fan power and room temperature does not cross the original temperature setpoint. Case 3: no overshoot in fan power and room temperature crosses the original temperature setpoint.

Lin et al. (2017) assumed overshoot in fan power was needed to return the building room temperatures back to their original setpoint. They found that the average room temperature tended to follow that of case 1 in Figure 1 and at the end of the event was close to the setpoint in the second half of the event, e.g., UP-DOWN events end with a higher average room temperature than the original temperature setpoint, and thus more cooling is needed to bring the room temperature back to the original temperature setpoint. Later, Raman and Barooah (2020) expanded the room temperature explanation to cases without overshoot, explaining that if the room temperature followed case 2 in Figure 1 and was closer to the setpoint in the first half of the event at the end of the event, then no fan power overshoot would be needed to bring the room temperature back to the original setpoint, e.g., UP-DOWN events would end with a lower average room temperature than the original setpoint, and thus less cooling is needed to bring the average room temperature back to the original temperature setpoint. From this explanation, the presence of overshoot, and whether energy is over- or under-consumed during the event is determined by if the average room temperature crosses the original temperature setpoint during the second half of the event or not.

However, an explanation for excess energy consumption that solely relies on room temperature may not be possible, as building physics is complex. For example, in the experimental work by Keskar et al. (2022), the fan power tended to follow case 2 (no overshoot), but the average room temperature follows case 1 (crosses the original temperature setpoint during the event). We denote this as case 3 in Figure 1; it is the typical case we observe in our experimental data. The explanation of energy consumption provided by Raman and Barooah (2020) based on average room temperature does not align with case 3. In the UP-DOWN case in Figure 1, after the event concludes, there is an under-consumption of fan power, yet the room temperature decreases. By only considering changes in room temperature it is expected that under-consumption of fan power would lead to the building temperature increasing above the original temperature setpoint, not decreasing. Similarly, in the DOWN-UP case, after the event concludes there is an over-consumption of fan power, yet the room temperature increases.

The past arguments for using room temperature to explain excess energy consumption are rooted in the assumption that fan power is directly related to airflow (e.g., through a cubic model), which may not always be the case in practice (Lin et al. 2024a). Under this assumption, changes in fan power would create prescribed changes in airflow and, subsequently, prescribed changes in building cooling service. In the context of sub-hourly HVAC load shifting, there has been no discussion of how changes in airflow (i.e., ventilation) could affect or explain excess energy consumption. Specific controllers, such as those designed by Wang et al. (2021), target control of the fan airflow directly. Other controllers have been designed to control duct pressure instead (Maasoumy et al. 2013). And still others have been designed to control fan speed (Hao et al. 2014). However, these studies do not mention excess energy consumption of the fans during load shifting, and there is still an unclear relationship between the airflow, pressure, and excess energy consumption. Further, direct control of the variables within the fan loop may produce unintended side effects, whereas GTA does not modify the fan control loop and, as such, may be more easily implemented in some buildings.

Overall, it is unclear what building services or variables compensate for the changes induced by load shifting and what factors contribute to these changes. Without understanding where excess energy is used within the building (or lost from the building) and what factors influence the building response, we cannot fully understand how the building is affected by load shifting or how the power consumption of the building will be changed by an event. Understanding the key factors that influence the buildings response and how building services are changed is critical for planning and deployment of grid-interactive buildings.

Methods

We used an iterative data-first method of analysis, in which we leverage experimental data to create informed hypotheses about building operation. This method is shown in Figure 2. We examine the experimental data to form hypotheses about potential sources of excess energy consumption caused by load shifting. Often, the experimental data was not enough to draw convincing conclusions. In these cases, we turned to simulation-based modeling to examine the dynamic states of the building. Our goal was not to model particular buildings accurately, but instead to use a generic VAV reheat building model to explore typical phenomena. Generally, in doing so, we would discover unexpected results, driving new hypotheses, which we compared against the experimental data. Building HVAC systems have many layers of control that interact through complicated physics. This iterative examination of data to create informed hypotheses allowed us to peel back these layers and garner a holistic understanding of building HVAC system response to load shifting.

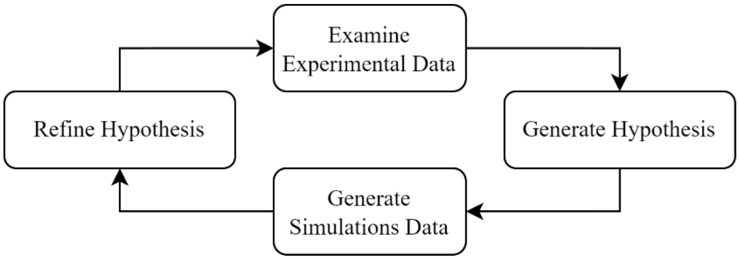

Figure 2: Flowchart demonstrating our data analysis method. This process was iterative as we continually formed and analyzed hypotheses based on both the simulation and experimental data.

All the events examined in this paper use open-loop GTA for fan power control. GTA indirectly adjusts VAV boxes' damper positions and the AHU fans respond to compensate for pressure changes. While the intent of GTA is to change the temperature in every room, in practice some rooms either have no response to the setpoint change or do not receive it (Lin et al. 2024a).

In past work, we have found that perceived excess energy consumption from load shifting may be influenced by poor testing design, i.e., selection of temperature setpoint changes that do not achieve perfect load shifting (Lin et al. 2024b). For these events, the excess energy consumption could be a byproduct of the setpoints selected and not the energy required to bring the building back to normal operation. We define energy-neutral events as load shifting events in which the buildings consume (approximately) the same amount of energy as the baseline energy consumption during the period of setpoint changes. By examining only those events that are energy neutral, we study how the building recovers from load shifting.

In the remainder of this section, we describe how the experimental data was acquired, giving context to the real building experiments conducted. We then describe the simulations that were performed for this paper to benchmark the experimental results. We discuss the filtering method that has been previously developed to select energy-neutral events. In the following section, we will discuss the results of our data analysis.

Experimental Data

This paper makes use of the SHIFDR dataset capturing fan submetering data, building automation system (BAS) data, and whole-building electric load data from sub-hourly load shifting experiments across 14 buildings and 5 years. Here, we use the preprocessed data from 2021, which includes 1-minute interval fan power and BAS control and monitoring data from six buildings located in southeast Michigan (Lin et al. 2023). All load shifting events in 2021 were 1-hour UP-DOWN or DOWN-UP events. This subset of data was selected as it contains more energy-neutral events than other years (Lin et al. 2024b). It also contains many baseline days in which no events occurred; these can be used as a benchmark for normal building operation. We note that each building in the SHIFDR dataset was given a pseudonym based on large lakes.

The six buildings range in construction year from 1941 to 2010 and in size from about 60,000 ft^2 to 280,000 ft^2. All buildings use VAV terminal reheat systems and are controlled with Siemens BASs. A full description of VAV terminal reheat systems can be found in Janis and Tao (2018); a brief summary is provided here. A chiller cools the supply air and a supply fan creates positive pressure in the ducts. VAV boxes connected to the supply duct control the amount of cool air that flows into each room or zone, regulating the room temperature. For the buildings considered in this study, the chiller is off-site (supply air cooling is via chilled water loops serving multiple buildings) and its power was not directly measured, as we only considered the fan power. Additionally, all of the buildings have return fans which help move air out of rooms via return ducts. Supply fans are controlled to maintain duct pressure, where the duct pressure setpoint is either constant or determined through a pressure reset method, adjusting the duct pressure setpoint depending on the position of the VAV box dampers. The return fans for these six buildings are controlled to match the return airflow and the supply fan airflow, minus some discount for exfiltration. For a detailed description of the buildings used in this study, please refer to the dataset (Lin et al. 2023) or companion paper (Lin et al. 2024a).

For the experimental data, the baseline energy consumption is unknown. However, several methods have been developed to estimate the baseline power. For this work, we employ a linear baseline method, in which we take a linear regression of the fan power data from 5 minutes before the event and 5 minutes after a 1-hour settling period (Keskar et al. 2020). This method assumes the fan power would have followed a linear path from the start of the event to after the event has settled. Lei et al. (2020) have demonstrated accuracy of this method for baselining fan power, which does not generally follow the same pattern as chiller power or whole-building electric load.

Simulation Model and Data

Simulations were generated using the Modelica Buildings Library version 8 (Wetter et al. 2014) and Dymola. The ASHRAE2006 five zone VAV reheat example model building was used for all simulations. The general control of this building is similar to that of the experimental buildings; a supply fan moves cooled supply air into a duct where VAV box dampers adjust the airflow into each room. Supply air is cooled by chilled water loops from an external source. The supply fan is controlled to maintain duct pressure, where the duct pressure setpoint is adjusted such that the maximum VAV damper position is 90%. There is no return fan and air leaves through the return duct due to positive room pressure. This model uses TMY3 weather data for the Chicago O'Hare Airport, which includes data for the entire year of 2006. A custom Modelica block was inserted into the model to modify the temperature setpoints of the rooms to simulate load shifting events, each day from 1 pm to 2 pm starting on August 1, 2006 and ending on August 11, 2006. Due to slow initialization of the model, the simulations were started on July 25, 2006, one week prior to the first event. Each simulation run yielded 11 one-hour long load shifting events each occurring on a different day.

Simulation-based testing was designed to generate energy-neutral events. Through a trial and error process, we determined sets of temperature setpoint changes that produced approximately energy-neutral events. Specifically, for both UP-DOWN and DOWN-UP events, the temperature setpoint increase was always set to 1°C. The temperature setpoint decreases for

UP-DOWN events were set to 0.25°C, 0.30°C, and 0.35°C, and for DOWN-UP events were set to 0.25°C, 0.27°C, 0.30°C, 0.35°C, and 0.37°C. For each set of temperature setpoint changes, we performed a full simulation, generating 11 events. The simulation model was numerically integrated using a DASSL solver with a variable step size. All data from the simulations were linearly interpolated to 1-minute resolution, to match the time resolution of the data in the SHIFDR dataset. The simulation baseline was generated by simulating the model building without temperature setpoint changes. This provides the exact baseline; no estimation is needed for baseline power.

Energy-Neutral Event Filtering

In ideal load shifting, the additional energy consumed during the power increase exactly matches the energy deficit during the power decrease, such that the total energy consumed during the event matches the counterfactual energy consumption if no event had occurred, i.e., the baseline energy consumption. We believe that some prior work has presented skewed HVAC load shifting efficiency statistics by including events that had significant changes in energy consumption during the event itself (not just during the settling period), i.e., the event did not achieve load shifting (Lin et al. 2024b). Therefore, we filtered our experimental and simulation results to include only events that are close to energy neutral.

We filter for energy-neutral events using the methods developed by Lin et al. (2024b). First, the linear baselining method is applied to non-testing days. Then, the energy consumption of the building during fictitious event windows on the non-testing days is compared to baseline energy consumption to find the average energy deviation associated with the linear baselining method, i.e., the average error in the baseline in terms of energy. Then, the linear baselining method is applied to the events and the excess energy consumption during each event is calculated. Events are energy neutral if the excess energy consumption during the event is within the average energy deviation associated with the linear baselining method. Additionally, events which do not result in a sufficient fan power response are excluded from the list of energy-neutral events.

We filter the simulation results using the same method. We found heuristically that applying the linear baselining method to the simulation result gives a reasonable tolerance for filtering energy-neutral simulation events. This is the only time we apply the linear baselining method to the simulated results.

Results

From our analysis we have developed several hypotheses attempting to explain why buildings consume excess energy after load shifting. Our analysis suggests that several routine building control mechanisms could be inadvertently changing the building's operating state and causing excess energy consumption. In this section, we discuss several of our hypotheses for the source of the excess energy consumption and provide supporting evidence, and then we provide a summary of suggested future experimentation needed to determine the source of the excess energy consumption and to explore how the excess energy consumption may be reduced.

Potential Sources of Inefficiency

We identify several explanations for the excess energy consumption of HVAC fans after short timescale load shifting. Unlike past work, we specifically analyze changes to the building ventilation system, not just changes to room temperature. We discuss in detail two causes related to the room and duct pressure: interference by the pressure reset controller and pressure sensor drift. In the case of pressure reset control interference, we show how excess energy consumption can be reduced by removing the pressure reset control. We then discuss other potential contributors to excess energy consumption.

Pressure reset control interference. To demonstrate how normal building control systems may behave undesirably during load shifting, we examine the pressure reset controller during the load shifting event. We hypothesize that some of the excess energy consumption during load shifting is caused by the pressure reset controller recovering its original duct pressure setpoint value at the end of the event. During this time, the fan requires more (or less) power to move air at a higher (or lower) pressure, and the excess energy consumption is the energy required by the fan to operate at a higher (or lower) duct pressure. We believe that this is the primary source of excess energy consumption in the ASHRAE2006 example building.

As duct pressure increases, the power required to move the same volume of air increases. The fan imparts more energy into the air to increase the air's pressure, from the ambient pressure to the duct pressure, when the duct pressure is higher. Similarly, when the duct pressure is lower, the fan imparts less energy into the supply air. A fan's cooling ability comes from delivering cool air to the room, regardless of the duct pressure. The static pressure reset controller aims to reduce the fan power by reducing the duct pressure to a minimum value that still satisfies the airflow constraints of the VAV boxes. For the ASHRAE2006 example building, the duct pressure setpoint is controlled via a PI controller that maintains the maximum VAV box damper to 90%. Within the experimental buildings, a similar control law is present, where the pressure level is changed if three or more VAV dampers are open more than 95% for 10 minutes.

We show results for energy-neutral simulation events in Figure 3, which shows the fan power, average room temperature across the 5 zones, supply airflow, and duct pressure. We show each event in blue and the average (at each time) in black. The data has been baselined; values represent the change from the baseline simulation in which no events occurred. The green and red lines indicate the start and end of the event, respectively. After the event, the pressure value takes the longest to return to normal, indicating the pressure reset loop is the slowest acting of the control loops.

We observe pressure setpoint changes in Figure 3 as changes in the duct pressure. Increasing airflow results from dampers opening, which triggers an increase in the duct pressure from the pressure reset controller, to maintain the maximum position of all the dampers at 90%. Similarly decreasing airflow results from dampers closing, which triggers a decrease in duct pressure from the pressure reset controller. In the UP-DOWN case, the pressure is lower than the baseline pressure during the settling period. Despite the fan power being lower than baseline power during settling, we see higher than baseline airflows. With more cold supply air, the room temperature expectedly decreases. The reverse is the case for DOWN-UP events, where the pressure during the settling period is higher than the baseline pressure, so there is less airflow than the baseline airflow despite the fan power being higher than the baseline fan power during the settling period. This relationship is weaker as seen by the smaller overshoot in supply airflow during the settling period for this event type.

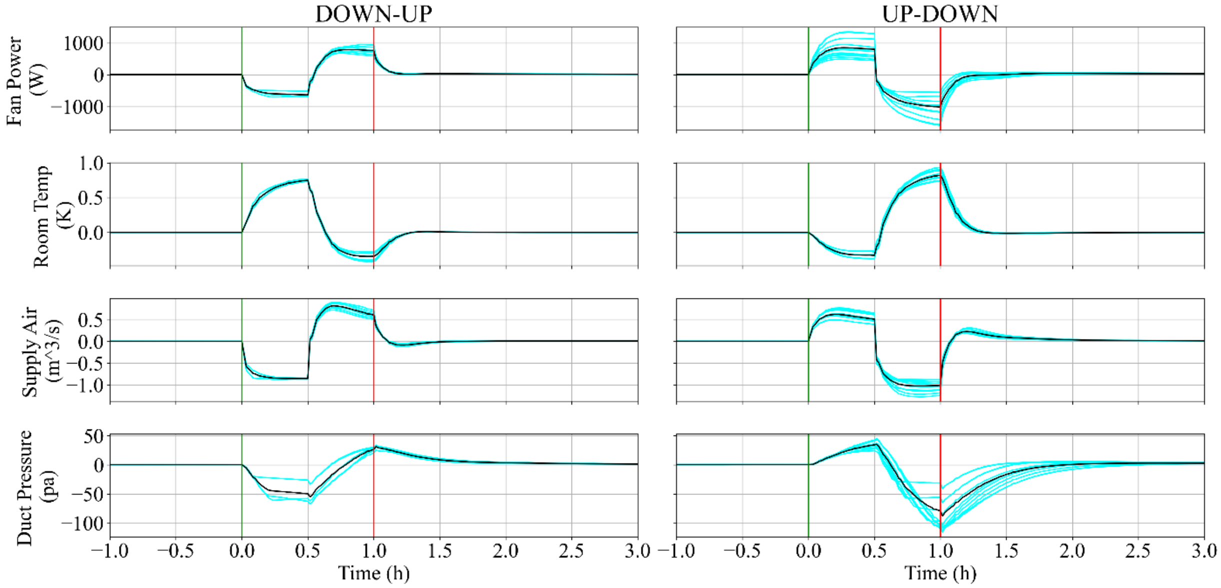

Figure 3: Fan power, room temperature, supply airflow, and duct pressure measurements from the simulations. Only energy-neutral events are shown. Individual events are shown in blue while the average (at each time) is shown in black. The green and red lines indicate the event start and end points. All events have been baselined, such that values reflect the change in values as compared to if no event had occurred.

Excess energy consumption caused by the pressure reset controller is not inherent to the building physics. In Table 1 we provide simulation results that show the impact of eliminating the pressure reset controller. We used a value of 150 pa for the constant duct pressure setpoint. Energy-neutral constant pressure simulation events were generated by repeating all simulations without the pressure reset controller and then applying the energy-neutral event filtering method. We compute the average excess energy consumption of the UP-DOWN and DOWN-UP events across all energy-neutral events for both control types by subtracting the baseline energy consumption from the building energy consumption between the start of the event and the end of a 1-hour settling period, and averaging the results across all energy-neutral events of that type.

Table 1. Average excess energy consumption of energy-neutral events for pressure reset and constant pressure simulations. For reference, the fan consumes an average of 2683.5 Wh in the baseline case over the event and settling periods.

Control type	Pressure Reset	Constant Pressure
DOWN-UP	88.06 Wh	67.29 Wh
UP-DOWN	-77.96 Wh	39.65 Wh

As shown in Table 1, by removing the pressure reset controller, the magnitude of the excess energy consumption for both event types becomes smaller, indicating it is closer to an ideal load shift. We note that while UP-DOWN events had negative excess energy consumption in the pressure reset case, they have positive excess energy consumption in the constant pressure case. Negative excess energy consumption indicates that the fans consume less energy than they would have otherwise. This shows that the building control design has an impact on the building's energy response, and disabling the pressure reset control during load shifting could reduce excess energy consumption.

Unfortunately, duct pressure was not measured in the SHIFDR dataset, and the presence of duct pressure resets is unknown. We suggest future experimentation to override the pressure reset controller to determine if a similar phenomenon to that observed in Table 1 is present in real buildings.

Pressure sensor drift. We hypothesize that fast airflow changes have unintended interactions with the pressure sensors in the building, causing slower settling than intended. We observed, in both the experimental and simulation data, a mismatch between supply and return fan airflow during and after an event. We refer to this as mismatch airflow and define it as the difference in supply airflow and return airflow measurements at any given time point. In Figure 4, we plot the fan power, along with the supply, return, and mismatch airflow for all energy-neutral DOWN-UP events from the experiments (first six columns) and simulations (last column). We normalized all data points by dividing by the average value during the plotted window so that data from different buildings can be directly compared. No baselining methods are applied to the data presented in Figure 4. We show each event in blue and the average (at each time) in black. The UP-DOWN events are not plotted but follow similar (but mirrored) trends to those in Figure 4.

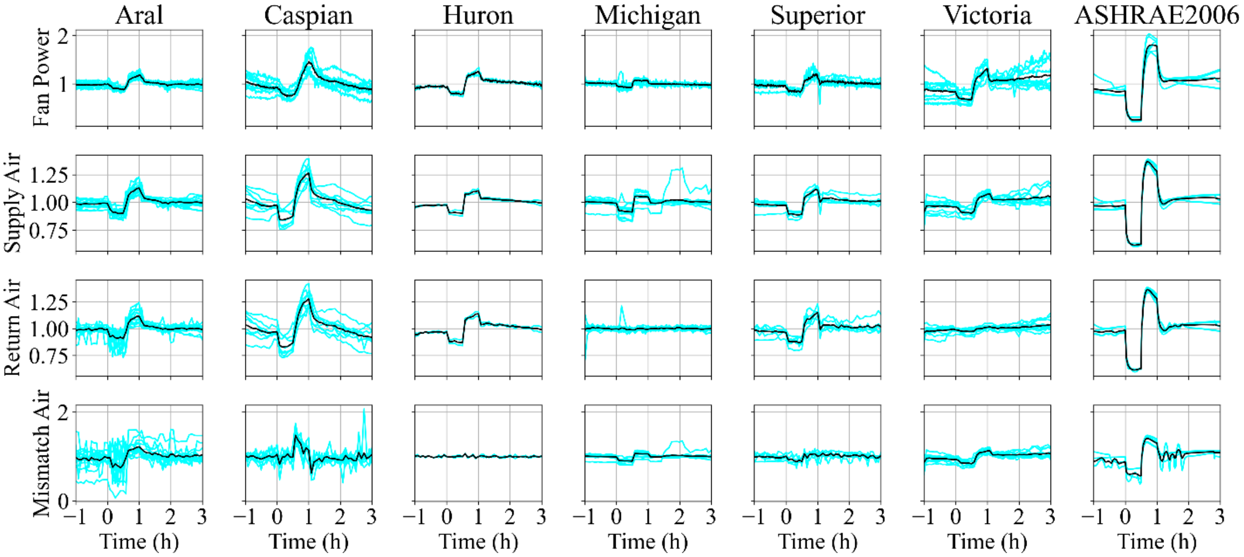

Figure 4. Fan power, supply airflow, return airflow, and mismatch airflow for energy-neutral DOWN-UP events. Individual energy-neutral events are plotted in blue, with the average (at each time) plotted in black. Values have been normalized such that the average across the plotted window is unity.

Figure 4 shows that the mismatch airflow changes during events, which we believe is caused by pressure changes in the rooms and the return fan controller taking time to match the supply fan airflow. In the experimental buildings, the return fans are controlled to match the supply fan airflow minus a discount for exfiltration. The sensors used to regulate duct pressure measure the relative pressure between the duct and the room where the sensor is mounted. If the mismatch airflow fluctuations significantly change the room pressure, the pressure sensor measurement may be impacted, i.e., the room pressure and, subsequently, the duct pressure measurement (which is relative to the room pressure) will drift from the ambient pressure. For example, increasing supply air without increasing return air would increase room pressure, making the duct pressure sensor measurement lower, triggering an increase in fan power. In the short term, while the return fan airflow is changing, this might cause slower than expected settling.

We may be observing the consequences of mismatch airflow affecting the room pressure in Figure 4. For the buildings Aral, Huron, and Superior, after each temperature setpoint change, there are two distinct patterns in fan power behavior: a large initial step followed by a slower drift in the direction of the step. We postulate that the large initial action is the response of the fans to the damper movement, while the slower drift is caused by room pressure changes while the supply and return fan controllers work to find an equilibrium. These building behaviors contrast with that of the building Michigan, where we observe large steps in fan power without any secondary action. We hypothesize that the lack of return fan response in the building Michigan plays a role in how pressure sensors are affected by the mismatch airflow.

Unfortunately, these possible phenomena cannot be explored with the ASHRAE2006 example building as the model measures duct pressure relative to the constant outdoor ambient pressure. Further, this model lacks a return fan. We suggest that future experimentation explore the role of mismatch airflow by repeating events with different control designs and/or tuning parameters for the return fans. There may exist a return fan control design that reduces the excess energy consumption of the events.

Other potential contributors to excess energy consumption. We have explored several additional hypotheses for the excess energy consumption; however, we lack the specific experimental data needed to test these hypotheses. In this subsection, we briefly describe these hypotheses, the relevant data that we are missing, and provide a brief discussion of whether the excess energy consumption could be reduced by changing the building control design.

We hypothesize that excess energy consumption of the fans could correlate to increasing fresh air intake during the event and/or settling window. Initial simulation results found this could be true; however, subsequent simulations did not show a clear trend. Additionally, SHIFDR does not include economizer airflows or control setpoints. Still, we believe that experimentation on the economizer function during load shifting would provide insight into how the building's ventilation services are affected by load shifting. We postulate that, if ventilation is a factor affecting excess energy consumption, changing the economizer control setpoints or control design during an event may reduce the excess energy consumption and be advantageous for load shifting.

We can also not rule out that the sensors in the building or the HVAC equipment are behaving erroneously, and the intended control function is not achieved. The excess energy consumption could be a byproduct of faulty control behavior. It is still unclear whether existing BAS infrastructure is sufficient to handle the control needed for DR (Vindel et al. 2023). Additionally, as discussed by Lin et al. (2024a), there are limitations on the sensor communication networks and network congestion can lead to sensor values not being reported in real time. The BAS data collected for SHIFDR was limited because of concerns about network congestion. Additionally, equipment faults may go undetected. For example, in the building Michigan, return fans did not respond to changes in supply airflow. The extent of sensor and equipment errors in the building are unknown without further monitoring and the extent to which they impact excess energy consumption is unclear. We suggest that future experimentation consider this, especially when designing closed loop control around building variables.

Another potential source of perceived excess energy consumption is the baseline modeling error. Excess energy consumption is sensitive to errors in the baseline (Lin et al. 2017). It could be the case that the perceived excess energy consumption is, in part, due to error in the baseline as opposed to being caused by the load shifting events. Clarity on what building states are modified during the event would help quantify the extent of baseline error.

A final possible explanation is that changes in average room temperature do account for the excess energy consumption, as described by Raman and Barooah (2020). Temperature sensors could provide erroneous measurements that do not accurately represent the room temperature. Poor mixing in the rooms could lead to the average room temperature not being represented by the temperature sensor reading. Additionally, not every room in the experimental buildings was monitored or responded to the setpoint changes (Lin et. al. 2024a), and rooms that were unmonitored could account for the expected room temperature changes. The extent of all of this is unclear. Resolving this would require collecting more room temperature data and, for the case of poor mixing, adding multiple temperature sensors per room. However, if the errors in average room temperature measurements are known ahead of time, the excess energy consumption can be eliminated by progressively adjusting the temperature setpoints used during load shifting events to account for the errors.

Roadmap for Future Experimentation

The results from our data analysis are incomplete. We have made several observations about how building states are changed during load shifting and how these changes might affect HVAC fan excess energy consumption due to load shifting. However, the SHIFDR dataset did not capture data related to ventilation and thus our findings are speculative. Without these key data points, the extent to which the excess energy consumption can be reduced remains unclear. Table 2 provides a summary of our hypotheses. For each hypothesis, we identify the data that is currently missing from the experimental dataset that prevents further investigation. We discuss whether the source of excess energy consumption can be reduced or not based on our understanding of the building control design. Finally, we propose future experimentation that should be conducted to determine the extent of the excess energy consumption from that source and whether excess energy consumption could be reduced through changes to building control design. Therefore, Table 2 provides an experimental roadmap for the future.

Table 2. Experimental Roadmap: Summary of hypotheses discussed in this paper, the data needed to further investigate each, the potential ability to reduce the excess energy consumption, and recommended future experimentation.

Hypothesis	Missing Data	Could it reduce excess energy consumption?	Future Experiments
Pressure reset control interference	Pressure setpoints, duct pressure	Yes, through pressure reset controller	Changing pressure reset controller
Pressure sensor drift	Absolute duct and room pressure, fan speed control signal	Yes, through return fan controller	Changing return fan controller
Changes in intake air	Intake airflow, economizer airflows and control measurements	Unknown	Changing economizer function
Building sensors or equipment are behaving erroneously	Independent data sensors	Unknown	Using independent (not BAS) sensors, testing on buildings without data congestions
Baseline modeling error	Unknown	Unknown	More accurate baseline methods
Average room temperature accounts for missing energy	Temperature from all rooms, more temperature sensors	Unknown	Applying progressive setpoint changes

Conclusion

In this paper, we provided a new analysis of the excess energy consumption by fans in VAV HVAC systems during sub-hourly load shifting DR events. Whereas previous work explained the changes in energy consumption through changes in room temperature, we focused our analysis on changes in ventilation. We found that during load shifting, several building control mechanisms may be behaving undesirably leading to excess energy consumption. We found that removing the pressure reset controller reduces the excess energy consumption, indicating that, at least in part, excess energy consumption is a byproduct of building control design. We identified several other building control mechanisms and factors that could contribute to actual or perceived excess energy consumption. If building control design is primarily to blame for excess energy consumption (versus inherent building physics), this is good news, as it indicates that excess energy consumption could be reduced by modifying the control during DR events. Finally, we outlined a roadmap of future experimentation that would allow us to determine the full extent of excess energy consumption and strategies that could be used to reduce it. This will be an important step for real-world implementation of short timescale DR, as it will help us identify the value and costs of DR, and design the mechanisms needed to properly incentivize it.

Acknowledgments

This work was supported by the National Science Foundation Graduate Research Fellowship Program under grant DGE 1841052. We are also immensely grateful to the facilities and operations staff at the University of Michigan for supporting our research by setting up experiments and providing insight into building operations. We thank Ian Hiskens, Jeremiah Johnson, Aditya Keskar, and Shunbo Lei for insightful discussions.

References

Afroz, Z., G.M. Shafiullah, T. Urmee, and G. Higgins. 2018. "Modeling techniques used in building HVAC control systems: A review." *Renewable and Sustainable Energy Reviews*, 83: 64-84.

Beil, I., I.A. Hiskens, and S. Backhaus. 2015. "Round-trip efficiency of fast demand response in a large commercial air conditioner." *Energy and Buildings* 97: 47–55.

Hao, H., Y. Lin, A.S. Kowli, P. Barooah and S. Meyn. 2014. "Ancillary Service to the Grid Through Control of Fans in Commercial Building HVAC Systems." *IEEE Transactions on Smart Grid* 5(4): 2066-2074

Janis, R.R. and W.K.Y. Tao. 2018. *Mechanical and Electrical Systems in Buildings, 6th edition* New York: Pearson

Keskar, A., D. Anderson, J.X. Johnson, I.A. Hiskens, and J.L. Mathieu. 2020. "Do commercial buildings become less efficient when they provide grid ancillary services?" *Energy Efficiency* 13(13):487–501.

Keskar, A., S. Lei, T. Webb, S. Nagy, I.A. Hiskens, J.L. Mathieu, and J.X. Johnson. 2022. "Assessing the performance of global thermostat adjustment in commercial buildings for load shifting demand response." *Environmental Research: Infrastructure and Sustainability* 2(1): 015003.

Lei, S., D. Hong, J.L. Mathieu, and I.A. Hiskens. 2020. "Baseline estimation of commercial building HVAC fan power using tensor completion." *Electric Power Systems Research* 189: 106624.

Lin, Y., J.L. Mathieu, J.X. Johnson, I.A. Hiskens, and S. Backhaus. 2017. "Explaining inefficiencies in commercial buildings providing power system ancillary services." *Energy and Buildings* 152: 216–226.

Lin, A.J., S. Lei, A. Keskar, I.A. Hiskens, J.X. Johnson, J.L. Mathieu, S. Afshari, D. Anderson, H. Li, S. DeMink, J. Dongmo, C. Flynn, P. Giessner, T. Kennedy, and K. Morgan. 2023. "SHIFDR Dataset - Michigan." May 2023. doi.org/10.7302/vmwh-ek70

Lin, A.J., S. Lei, A. Keskar, I. Hiskens, J.X. Johnson, and J.L. Mathieu. 2024a. "The Sub-Metered HVAC Implemented for Demand Response Dataset." *ASME Journal of Dynamic Systems, Measurement, and Control* 146(1): 011107.

Lin, A.J., S. Lei, J.L. Mathieu. 2024b. "Experimental Investigation of Building HVAC Load-Shifting Efficiency." *Power and Energy Society General Meeting*

Liu, J., L. Yu, R. Yin, M. Piette, M. Pritoni, A. Casillas, and P. Schwartz. 2023 "Factors Influencing Building Demand Flexibility." Lawrence Berkeley National Laboratory. doi.org/10.20357/B71882

Maasoumy, M., J. Ortiz, D. Culler, and A. Sangiovanni-Vincentelli. 2013. "Flexibility of Commercial Building HVAC Fan as Ancillary Service for Smart Grid." *Proceedings of Green Energy and Systems Conference.*

MacDonald, J.S., E. Vrettos, and D.S. Callaway. 2020. "A Critical Exploration of the Efficiency Impacts of Demand Response From HVAC in Commercial Buildings." *Proceedings of the IEEE* 108(9): 1623-1639

Raman, N.S. and P. Barooah. 2020. "On the Round-Trip Efficiency of an HVAC-Based Virtual Battery." *IEEE Transactions on Smart Grid* 11(1): 403-410.

Taylor, J.A., S.V. Dhople, and D. S. Callaway. 2016. "Power systems without fuel." *Renewable and Sustainable Energy Reviews* 57: 1322–1336.

Vindel, E., B. Akinci, and M. Bergés. 2023. "A critical investigation of the readiness of VAV systems to adopt grid-interactive capabilities." *Energy and Buildings* 286: 112974,

Vrettos, E., E.C. Kara, J. MacDonald, G. Andersson, and D.S. Callaway. 2018. "Experimental Demonstration of Frequency Regulation by Commercial Buildings—Part II: Results and Performance Evaluation." *IEEE Transactions on Smart Grid* 9(4): 3224-3234

Wang, J., S. Huang, D. Wu, and N. Lu. 2021. "Operating a Commercial Building HVAC Load as a Virtual Battery Through Airflow Control." *IEEE Transactions on Sustainable Energy* 12(1): 158-168.

Wetter, M., W. Zuo, T.S. Nouidui, and X. Pang. 2014. "Modelica Buildings Library." *Journal of Building Performance Simulation* 7(4): 253-270