Stay cool and be flexible: Energy-efficient grid services using commercial buildings HVAC systems

Aditya Keskar, North Carolina State University
Shunbo Lei, University of Michigan
Taylor Webb, North Carolina State University
Sarah Nagy, North Carolina State University
Han Lee, University of Michigan
Ian A. Hiskens, University of Michigan
Johanna L. Mathieu, University of Michigan
Jeremiah X. Johnson, North Carolina State University

ABSTRACT

The thermal inertia of commercial buildings allows us to shift their power consumption on minutely to hourly timescales to provide grid services while maintaining occupant comfort. Our experimental research has shown that buildings providing these services may consume more energy than they would under normal operation. In this paper, we explore this phenomenon by analyzing results from over one thousand experiments on eight campus buildings in Michigan and North Carolina. This builds upon our prior work that drew preliminary conclusions based on experimental results from three buildings in Michigan (presented at ACEEE 2018). In our experiments, we manipulate each building's thermostat setpoints using pre-defined setpoint signals that cause the building to shift its power consumption with respect to its baseline. Expanding on our prior research using square-symmetric setpoint signals, we implemented a variety of new signals (e.g., ramping the setpoints) that enable an enhanced understanding of the physics and control response of HVAC systems. We analyze the fan power response and investigate impacts on the building chilled water and terminal reheat systems, allowing us to conduct a more holistic energy efficiency impact assessment. We also quantify the impact of load shifting on the cooling service provided to the building thermal zones. Our paper provides new insights that help inform the design of building response strategies that mitigate inefficiency. This analysis also underpins the experimental design for the next phase of testing in Summer-Fall 2020.

Introduction

The electricity grid needs flexible resources that can maintain supply-demand balance across a wide range of time scales. As the penetration of renewable resources increases, achieving this balance will require increasing flexibility on the grid. This need is exacerbated as fossil-fuel based resources that previously provided such flexibility (e.g., existing thermal generators with high rotational inertia and load-following governor systems) are increasingly displaced. Efficiently leveraging new sources of flexibility to ensure reliable and economic power grid operation is an important step in achieving deep decarbonization of our energy system (Davis et al. 2018). In recent years, considerable research effort has been devoted to investigating the potential of commercial building heating, ventilation, and air conditioning (HVAC) systems for providing a variety of grid services (e.g., Watson et al. 2006; Zhao et al. 2013; Hao et al. 2014 and Cai et al. 2018). A recent report estimated the potential for 200 GW of cost-effective flexible load services that could be utilized on the U.S. grid by 2030 (Brattle Group, 2019).

Demand response has often been restricted to load shedding events, reducing the system load in times when the grid is operating near peak (typically on hot summer days). However, research has shown that demand response has tremendous potential for ancillary services like frequency regulation (e.g., Lin et al. 2015) and, as our research (and other prior works like Beil et al. 2015) has shown, for load shifting on minutely to hourly timescales. The timescales of the experiments presented in this paper are consistent with those of real-time energy markets, which occur every 5 to 15 minutes. The use of commercial buildings for a wide range of grid services across multiple timescales will require a much more complete understanding of building impacts. Viable control strategies should take into account the possibility of HVAC systems consuming more energy than they would under normal operation. This is illustrated by our previous work, presented at the 2018 ACEEE summer study (Keskar et al. 2018; Keskar et al. 2019), which showed that HVAC fans consume more energy when providing load shifting than under normal operation. The efficiency of building response and the additional energy consumption depended on the polarity of the setpoint signal given to the building thermostats (i.e., whether the setpoint temperature was first increased or decreased). These results were consistent with prior experimental work conducted at Los Alamos National Laboratory (Beil et al. 2015) but inconsistent with previous modeling work (Lin et al. 2017). One of the recurring themes in this research domain has been the apparent mismatch between the experimental and modeling work when it comes to assessing the efficiency of building response. To gather an enhanced understanding of the physics and control dynamics of building response, we have conducted a wider-scale experimental investigation into this phenomenon.

Our new work, reported here, broadens the scope of our research, while providing further insights into some of our prior assumptions. In 2019, we conducted over one thousand experiments on twenty campus buildings at North Carolina State University (NCSU) and the University of Michigan (UM). The supply and return fans in the buildings were instrumented. We then manipulated the building temperature setpoints using pre-defined signals (square, ramp, and successive, as shown in *Figure 1*) that changed the fan power consumption of the building with respect to its baseline. We also collected building automation system (BAS) data that allowed us to quantify the response of the chilled water and terminal reheat systems. This data provided new insights into the response of other grid-connected HVAC subcomponents when we use fans for shifting electricity demand on minutely to hourly timescales. We previously assumed the change in room temperatures would be negligible due to the vast thermal inertia of commercial buildings. This assumption is revisited by quantifying the *effective cooling service* (ECS) provided to a building due to load shifting. This is achieved by collecting a subset of room temperatures across different building thermal zones using the BAS. The results presented for the fan power, chiller, terminal reheat, and zone temperature response provide a holistic assessment of the impact that load shifting can have on commercial building HVAC systems.

In this paper, we present results from eight campus buildings in order to conduct a comprehensive assessment of the impact on the energy efficiency of response of commercial HVAC systems when used for load shifting. We first describe the experimental setup by providing a detailed description of the setpoints signals implemented, description of the buildings, instrumentation, data collection, and execution of experiments through the BAS. Then, we describe the metrics we used to quantify the efficiency of fan, chiller, terminal reheat response, as well as the method to quantify the effective cooling service provided to the building thermal zones. In the results section, we provide a detailed discussion of two buildings and summarize the results of the other buildings. We conclude by providing our key findings and highlight the future direction of our research.

Experiments conducted

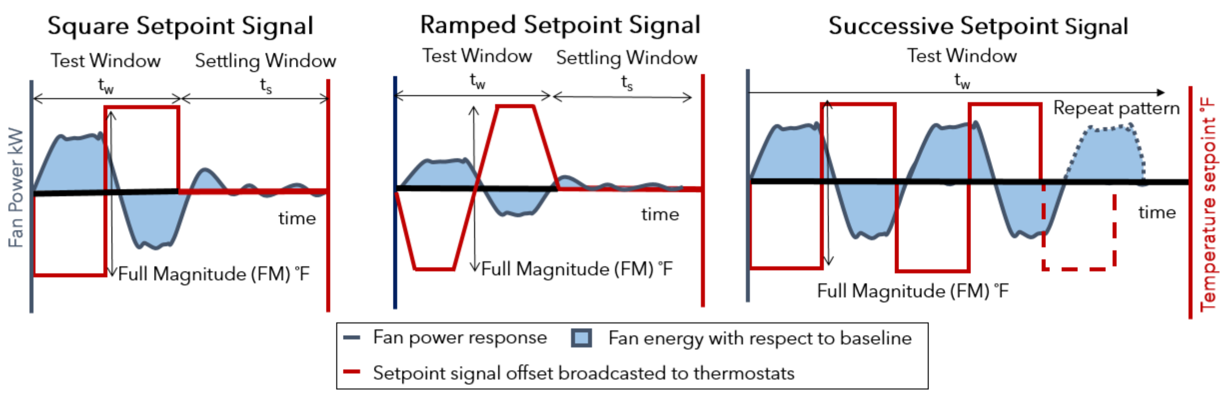

Figure 1. Thermostat setpoint shapes implemented on the eight buildings and the resulting aggregate fan power response relative to a baseline. An Up (Down) setpoint change leads to a Down (Up) fan power consumption change.

As shown in Figure 1, we analyze three types of setpoint signals given to the buildings. The setpoints are broadcast to the building thermostats within testing window t_w. We also analyze the response of the building after the test, over a pre-defined settling window t_s during which the building settles back to baseline (unperturbed setpoint) operation. The total window of analysis is the sum of the testing and settling windows ($t_t = t_w + t_s$). Our previous work found that the polarity of signals given to the building had a significant impact on the efficiency of building response. Specifically, building response was more efficient when the power consumed by the building was increased and then decreased with respect to its baseline (Up-Down power) than when the power consumption was first decreased then increased (Down-Up power). The efficiency of building response was quantified by the round-trip efficiency (RTE, a standard metric to quantify efficiency of energy storage) and the additional energy consumption (AEC, the additional kWh consumed by the building over its baseline). However, in this paper, we use only the AEC since our previous work found that it is more robust than the RTE for quantifying the efficiency of building response. To allow for direct comparison with our previous results (Keskar et al. 2019), we conducted 1-hour square events (Up-Down and Down-Up) at each of the eight buildings, as shown in Figure 1 (left). We also designed two new tests to investigate new hypotheses that enable an enhanced understanding of the response of the HVAC system.

The first new test, as shown in Figure 1 (middle) was a ramped setpoint test to investigate if ramping the setpoint signal could improve the efficiency of fan power response. In our previous work, one of the hypotheses we had proposed for the inefficiency of the Down-Up tests was the large setpoint step up applied in the middle of the test which causes the aggregate fan power consumption to rebound over its baseline.

The second new test, as shown in Figure 1 (right), were successive tests which were inspired by the modeling work conducted at the University of Florida (Raman and Barooah 2018a, and the follow-up work Raman and Barooah 2018b). Their analytical and simulation work found that the round-trip efficiency converges to unity as a building is subjected to successive load shifting. In other words, the building response is 100% efficient when it is repeatedly perturbed to provide load shifting grid services. Their model suggested that the indoor temperature deviation is not zero mean (building gets slightly warmer). In their follow-up work (Raman and Barooah 2018b), the authors constrained the average temperature deviation in their model to zero which caused the RTE to converge to values less than one since there was additional energy consumed to maintain the temperature constraint (i.e., not allowing the warming to occur). To the best of our knowledge, no prior experimental work has investigated the impact on building efficiency caused by successive load shifting.

Experimental setup

Building selection and characterization

Table 1. Building parameters

Building name	Eng. Bldg. II	Eng. Bldg. III	Park Shops	Sullivan Shops III	Bob & Betty Beyster Bldg.	Thayer Bldg.	North Quad. Complex	Dana Bldg.
Location	NC	NC	NC	NC	MI	MI	MI	MI
Year built	2005	2010	1914	2011	2005	2006	2010	1901
Building type	Lab/office	Lab/office	Office	Office	Classroom /office	Office	Classroom /office	Classroom /office
Area (ft^2)	202,400	175,000	50,000	13,500	104,100	59,800	288,400	117,100
Annual energy consumption (MWh)	4,284	4,649	761	Unavail.	3,160	508	3,979	1,595
HVAC system type	Chilled water based	Chilled water based	Chilled water based	Direct Expansion	Chilled water based	Chilled water based	Chilled water based	Chilled water based
Chilled water source	Central plant	Central plant	Central plant	N/A	Central plant	Central plant	Central plant	Central plant
#Setpoints controlled	265	335	85	15	193	94	124	162
#AHUs	6	10	2	1	3	2	16	3
#Fans instrumented (#supply fans)	16(10)	14(7)	4(2)	1(1)	7(3)	2(1)	4(2)	3(2)
2019 testing period	April-October	April-October	May-October	June-October	June-September	August-September	July-September	July-September
Testing times	9-11 am, 1-3 pm	9-11 am, 1-3 pm	9-11 am, 1-3 pm	9-11 am, 1-3 pm	9-11 am, 1-3 pm	9-11 am, 1-3 pm	9-11 am, 1-3 pm	9-11 am, 1-3 pm
#Total tests conducted	183	155	131	99	52	32	32	40

The test buildings have multizone single duct variable air volume (VAV) systems with terminal reheat, with the exception of one building that has a direct expansion system (the refrigerant cools the air directly as opposed to in chilled water systems where the air is cooled using air to water heat exchangers). The chilled water is supplied either by a central campus chilled water loop or by a standalone chiller. The buildings were selected based on having a sufficient number of controllable setpoints available on the BAS, in addition to ensuring that we cover a wide variety of building sizes, types, HVAC system layouts, and number of fans, as shown in Table 1.

Instrumentation and BAS data collection

Figure 2. Current sensors installed on a single phase of a fan at Engineering Building-2 (left) and on a chiller at Toxicology Building (right) at NCSU.

We submetered 55 subcomponents (54 fans and 1 chiller compressor) at seven NCSU buildings and 90 subcomponents (89 fans and 1 chiller pump) across 13 UM buildings (Figure 2). To instrument the fans and chiller equipment, we installed current sensors (*Onset CTV-D 20-200A*) on a single phase of the supply fan and return fans belonging to the air handling units (AHUs). The current probes were attached to data loggers (*Onset- HOBO 4 Channel U120-006*). Using constant voltage and power factor assumptions, we estimated per-minute three-phase fan power consumption[1].

We also collected numerous building trends through the BAS. The following points were collected from each building when available: subset of zone temperatures and temperature setpoints**,** VAV damper positions, reheat valve positions, chilled water flow, chilled water supply temperature, chilled water return temperature, building steam flow, outside air temperature, outside air humidity, and electrical load data. Due to the large number of data points, some of these trends were only available at a lower temporal resolution (e.g., every five minutes). In this paper, we present the results of eight buildings, which offer the highest number of tests and most-complete datasets from the BAS.

---

[1]Ideally, we would have installed three-phase power meters on each of the fan VFDs. However, this was impractical due to the large number of fans and associated high cost. We would expect both the voltage and the power factor to vary throughout the day. The extent of the variation and impact on our results is a subject of current research.

Shifting building demand using Global Thermostat Reset

To shift the electrical load of the building, we utilize the global thermostat reset (GTR) methodology in which we offset the VAV temperature setpoints through the BAS. The change in VAV setpoints causes the VAVs to open/close the dampers to control the amount of air that is entering a thermal zone in order to maintain the newly commanded setpoint. (In our experiments, we control a majority of the VAV setpoints, only omitting building zones with sensitive loads such as lab equipment and hospital operating rooms, which require tight temperature control.) Through these actions, the increase/decrease in static duct pressure causes the fans to change their power consumption. Goddard et al. 2014 provides a comprehensive overview of other control strategies that can be used to trigger a change in commercial HVAC power consumption and an insightful discussion on the relative advantages offered by GTR.

At UM, the tests were coded into the BAS on the field panel of each individual building to control the VAV setpoints of that building. At NCSU, a Python program running on a cloud hosted virtual server was used to change the VAV temperature setpoints in all the buildings (i.e., each building receives the same change in setpoint command). It is not required for all buildings to receive the same setpoint command (but to ensure sufficient repeatability of tests we gave the same setpoint command to all buildings). A significant advantage provided by the approach at NCSU is scalability, as adding buildings is as simple as including the VAV network numbers in the central Python program. The program relies on an open-source library called Bacpypes (Bender, 2019), which enables the Python code to communicate using the BACnet/IP protocol (a communication protocol developed by ASHRAE for BASs). A disadvantage of this method, especially on a campus network, is that significant coordination with information technology organizations is required to ensure network firewall settings do not block the BACnet traffic from reaching the buildings. Another disadvantage of the approach at NCSU is the slower communication rate capability of the field devices compared to the virtual server that communicates the setpoint signals to them. This difference can overload field controllers with BACnet traffic and result in decreased controller performance. To overcome this challenge, we implemented a parameter in the Python program which limits the rate of communication to field devices. This is an important consideration when using this methodology for providing grid services in other buildings since different field controllers have different rates of allowable traffic.

Test parameters

Section A.1 of an Electronic Appendix (link provided at the end of the paper) contains a table that summarizes the tests we have analyzed in this paper. The table shows the full magnitude (FM, see Fig. 1) of the setpoint change and the duration of the testing and settling windows for each type of test: square, ramped, and successive. The ramped signals were only implemented at NCSU. We used two ramp rates for the tests: slow ramping which initially ramps the setpoint over 1/6 of the duration of t_w and a fast ramp that ramps the setpoint over 1/12 of the duration of t_w. To investigate the overall impact of ramping the setpoints on efficiency, we present the results of both the slow and fast ramped tests together (t_w and FM were kept consistent across both types of tests). Different FMs were used at UM and NCSU (4˚F vs. 6˚F) since initial experiments found different setpoint perturbations were needed to elicit a response clearly visible in the data.

Evaluating response of Fan, Chiller, Terminal Reheat, and Cooling Service

We have observed that sometimes the building HVAC system does not respond as expected. Therefore, we developed a method to remove outliers. The experiments were designed to produce a fan power response that is approximately symmetric during the testing window and visually observable from the data. Therefore, we removed tests that fail to meet symmetry (insufficient response in either the Up or Down direction) and magnitude thresholds (insufficient response with respect to the estimated baseline) using the outlier criteria listed in Keskar et al. 2019. The tolerance for symmetry, denoted ε_s, always takes the value 20% (e.g., if response is higher above the baseline, the kWh response below the baseline should be at least 20% of the response above the baseline in t_w), whereas the magnitude tolerance ε_m is assigned per building by inspecting the time-series data. Values for ε_m are provided in Section A.2 of the Electronic Appendix. In the future, we will also incorporate the change in mean damper positions into the outlier criteria since it is the opening/closing of dampers that causes the change in static duct pressure, which consequently causes the change in fan power consumption.

Chilled water systems are often the biggest drivers of energy consumption in commercial HVAC systems. Initially, we expected the chilled water system to be largely unresponsive due to its high time constant (Goddard et al. 2014). However, our previous work (Keskar et al. 2019) found the chilled water system to have a non-negligible response to the tests, though that result was based on data from one building. In this work, we conduct a more in-depth analysis using data from six buildings. Depending on the data available in the BAS, we use either the chilled water flow rate or the instantaneous tonnage of cooling used by the building. We determine the cooling load tonnage of a building by multiplying the difference between chilled water return temperature and supply temperature, the chilled water flow rate, and a constant that accounts for specific heat capacity (see Section A.3 of the Electronic Appendix). The required data is collected through the BAS at 1-minute or 5-minute intervals.

Terminal reheat systems have two main functions. One is to help control the humidity level in the thermal zone by reheating the air supplied to the zones and the other is to help maintain occupant comfort when the damper control loops are unable to maintain the desired zone setpoints. We investigated the impact of the experiments on these systems by collecting a subset of reheat valve position data. To ensure these data act as a comprehensive proxy for the impact on the entire reheat loop and subsequent analysis, we selected reheat valve data across rooms distributed throughout the buildings. We also collected a majority of room temperature and damper position data from the corresponding VAV boxes. The buildings that we investigated provided terminal reheat using hot water heat exchangers that were supplied from either a standalone boiler or a central campus steam loop. In prior experimental work (Goddard et al. 2014), the authors disabled the terminal reheat to ensure it would not engage. We were not able to disengage the reheat in our experiments.

Efficiency of fan power response

We use the AEC metric as defined in Keskar et al. 2019 (i.e., the energy consumed by the fan(s) above its baseline minus the energy consumed by the building below its baseline) to compute the fan power efficiency in the testing window t_w, settling window t_s, and total window

t_t. Since the tests are designed to achieve energy neutral load shifting, we would expect the AEC in t_w to be 0 kWh. However, due to non-symmetrical response in the Up and Down directions the metric takes a non-zero value. To compute the AEC, we fit a linear baseline using least squares to fit data 5 min before the beginning of t_w and 5 min after the end of t_s. Keskar et al. 2019 show the performance of this method to estimate the baseline by computing the error of the technique when used on baseline days. Lei et al. 2019 provide a comprehensive overview of a variety of different fan power baseline estimation methods. The two papers show that the linear baseline using least squares performs well in estimating baseline fan power.

Efficiency of chiller response

We compute the additional chiller consumption (ACC) with respect to the baseline to quantify the efficiency of chiller response. The building's baseline cooling load is estimated by a linear regression model similar to *Baseline Method-2* in Mathieu et al. 2010 used for whole building electric load prediction. Note that both whole building electric load and chiller cooling load are strongly correlated with outside air temperature. Here, we use 5-min interval cooling load data from the building chilled water system. Independent variables in the regression model include time of day, average cooling load from 7:00 am to 9:00 am, and the outside air temperature. Specifically, the nonlinear dependence of the cooling load on the temperature is captured by a piecewise linear approximation with six segments. The regression model is trained using data from the same building on days without tests, referred to as baseline days. For some buildings with insufficient baseline day data, we also include data from test days, but outside the experiment windows.

In predicting or forecasting the cooling load, calibration methods are commonly adopted for error reduction. For example, after a regression model produces the prediction of next time step, one can calibrate it by subtracting from it the regression model's average error in the previous two hours (Fan and Ding 2019). We adopt a similar calibration method here. In the total window t_t, the regression model's errors are unknown as the measured cooling load deviates from the baseline due to our tests. Still, before and after the total window, the regression model's errors are known and can be obtained by comparing the model prediction to the cooling load measurement, which is the true baseline. Let t be the index for time steps within the total window t_t, and $t_{\pm 1.5\text{hrs}}$ be the window covering 1.5 hours before and 1.5 hours after time t. Note that $t_{\pm 1.5\text{hrs}}$ is different for each time t. Let $t_{\pm 1.5\text{hrs}}$ - t_t represent the window containing time steps within $t_{\pm 1.5\text{hrs}}$ but outside t_t. In this work, we calibrate the regression model's baseline prediction at each time t by subtracting from it the average of the model's known errors within the window $t_{\pm 1.5\text{hrs}}$ - t_t. In the Electronic Appendix Section A.4, we quantify the accuracy and bias of our baseline method. We do this by computing the ACC at event times on baseline days (which should be zero if the baseline method is perfectly accurate) to obtain the average baseline error (bias) and standard deviation (accuracy). We also use this method to quantify the baseline error for the terminal reheat and effective cooling service.

Efficiency of terminal reheat response

To compute the terminal reheat power consumption, we would ideally collect data on the supply and return water temperature and flow rate, as well as the amount of air (in cfm) delivered by the VAV to the room. However, this data was not available from the BAS. Instead, we offer insights into the additional reheat consumption (ARC) by computing the change in mean reheat valve position (average of all collected reheat valve positions) in the testing window t_w and settling window t_s. The baseline method used is similar to the fan power method by fitting a linear least squares baseline using the mean valve position 5 min before the test and 5 min after the settling window. The ARC metric thus gives us a proxy for the impact on the terminal reheat systems and is quantified in units of %-hours. We present the error associated with this method in Section A.4 of the Electronic Appendix.

Effective cooling service provided

Changes to the level of cooling service provided to the building is of fundamental importance when investigating the efficiency of HVAC subcomponent response, as well as understanding the physics behind the inefficiency seen in our prior work for two reasons. One, it helps us understand the change in cooling provided across different types of events, enabling us to understand whether certain types of tests can lead to more or less cooling provided to the thermal zones. Two, understanding the change in cooling service provided helps us investigate the relation between the quality of service (cooling) and the cost of service (additional energy consumed by subcomponents). We use the average of a subset of zone temperature trends of the building to quantify the effective cooling service (ECS) provided to the thermal zones. This is calculated based on the deviation in temperature with respect to the baseline. Cooling setpoints in campus buildings were held at 72˚F or other values and thus room temperature deviations were minimal during baseline operation. We compute the baseline by fitting a linear least squares baseline using data 5 min before t_w and 5 min after t_s. The ECS is then quantified by subtracting the mean room temperature from its baseline and is computed in deg-hours units. A higher value of ECS indicates a larger temperature deviation above baseline, and consequent warming of thermal zones. Again, we present the error associated with estimating the baseline for the mean room temperature in Section A.4 of the Electronic Appendix.

Results

In this section, we detail the results from the experiments that were conducted across the eight campus buildings. Figure 3 shows the time-series plots for the three types of signals at Engineering Building-2. The figure shows how the different subcomponents of commercial HVAC systems react when we shift the HVAC fan power load. We see the change in setpoints (row one in Figure 3) causing the dampers (row five) to open/close to meet the new setpoint, which causes the aggregate fan power consumption (row 2) to decrease/increase to maintain the desired static duct pressure (for this building the values varied 0.8-1.2 in. wc depending on the AHU). The change in zone temperature (row one) and reheat valve positions (row four) show that their shifts with respect to their estimated baseline follow the polarity of the setpoint signal. We also see the response of the chilled water system (row 3) to the three different signals.

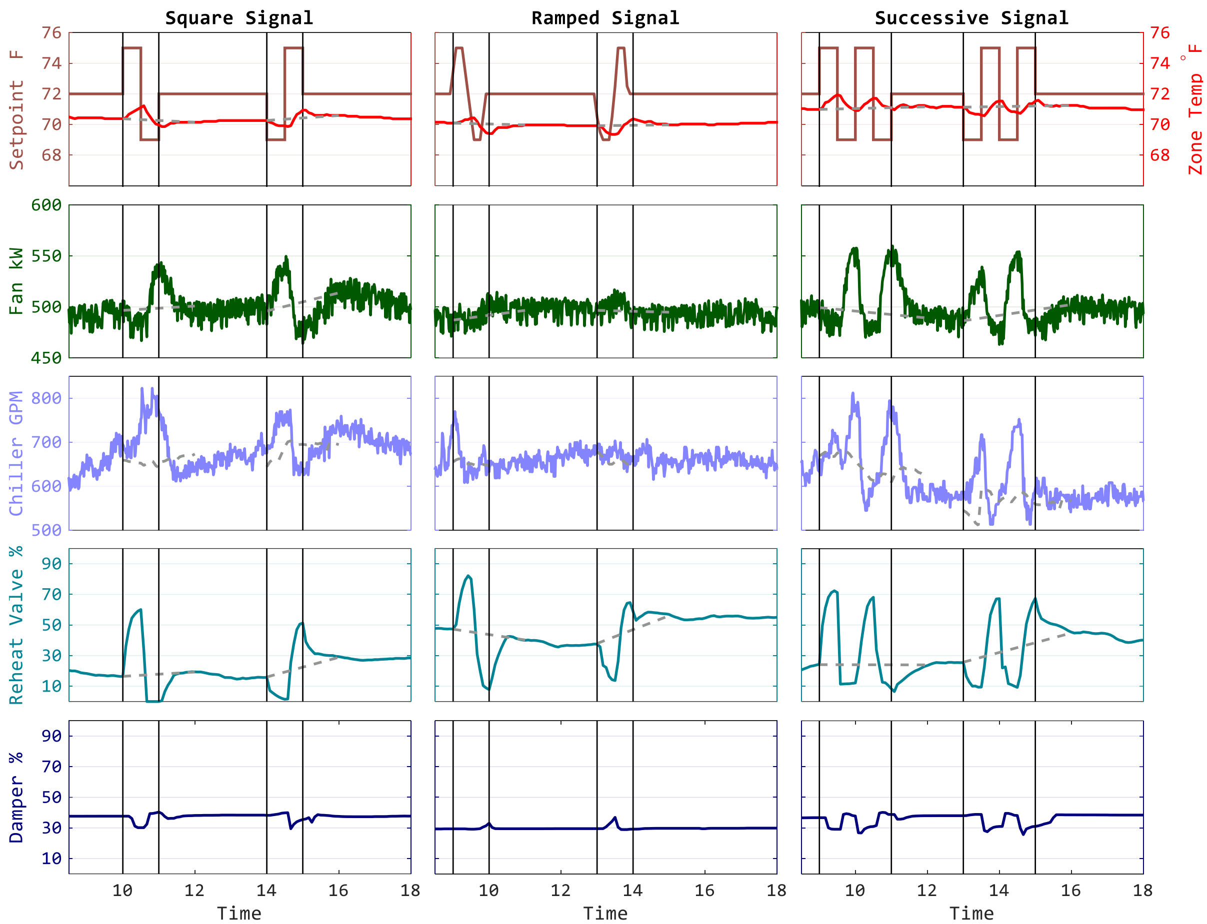

Figure 3: Representative results from Engineering Building-II for square, ramped, and successive events. The five rows (top to bottom) show the setpoint signals broadcasted (and corresponding changes in zone temperature), aggregate fan power consumption, chiller flow from campus chiller loop, average reheat valve positions, and average VAV damper positions, respectively. The grey dotted line shows the estimated baselines for the different subcomponents.

Below we present detailed results and analysis from two buildings, one at NCSU and the other at UM. Table 2 provides a summary of the different efficiency metrics computed for all the buildings.

Engineering Building-II

A total of 183 events were conducted at Engineering Building-II. As row 1 of Figure 4 (left) shows, the Down-Up square event lead to more fan power comsumption (higher AEC) on average in the total window t_t than the Up-Down square event, which is consistent with our prior experimental results. We also notice a corresponding increase in the amount of additional chilled water (increase in ACC in t_t) as seen in row 3 of Figure 4 (left), which seems to indicate that the Down-Up square signal also draws more power from the chilled water system. We also find that an increase/decrease in ARC values from t_w to t_s for successive and ramped events leads to a corresponding increase/decrease in ECS provided to the thermal zones (row 2 and row 4 of Figure 4). The non-zero ARC values in row 2 of Figure 4 are significant since we found the reheat to respond (as expected) with the the same polarity as the setpoint signal unlike the fans which respond with the opposite polarity. This indicates that system compensates for the change in cooling by using additional reheat power which is potentially one of the drivers behind the change in effective cooling service provided to the thermal zones. The median ECS values across all test types, as seen in row 4 of Figure 4 (~ 0°F-hours), indicate that the temperature deviation casued by the load shifting is minimal when we shift the fan power load. We find that ramping the setpoints significantly reduces the mean ACC compared to square signals, as seen in row 3 (middle) in Figure 4. We also find a decrease in AEC for Down-Up signals indicating an improvement in efficiency of response (row 1 (middle) in Figure 4). However, compared to square signals, ramped signals generally result in smaller reponses of HVAC subcomponents. The relationship among setpoint ramp rates, response maganitudes, and energy efficiency needs further investigation in future work. For the successive events, we find an overall increase in fan and chiller consumption (row 1 (right) and row 3 (right) in Figure 4).

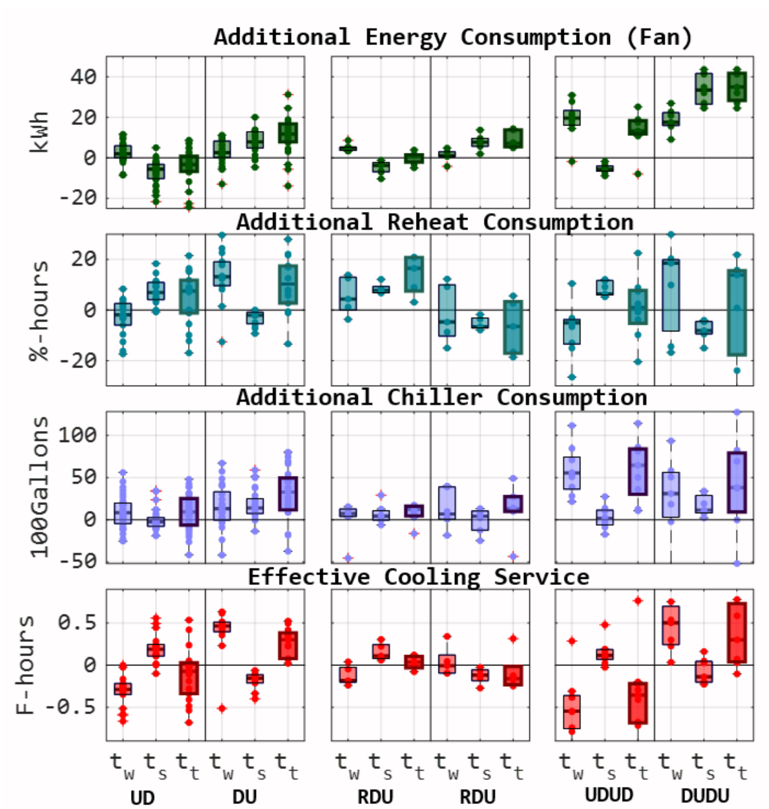

Figure 4: Efficiency metrics for fan, reheat, chiller and effective cooling service for the three different types of signals for Engineering Building-2. The dots show the respective metrics for one event. UD: Up-Down, DU:Down-Up, RUD: Ramped Up-Down, RDU: Ramped Down-Up, UDUD: Up-Down-Up-Down, and DUDU: Down-Up-DownUp. The edges of the boxplots showing the metrics for t_t have been darkened.

Thayer Building

A total of 32 events were conducted at Thayer Building. We find that the mean AEC of the square Up-Down is less than the mean AEC of Down-Up as seen in row 1 (left) of Figure 5, indicating a more efficient fan response. We also find the ACC for the square Down-Up event to consequently be higher than the Up-Down square events, as seen in row 3 (left) of Figure 5. The t_w for the square and successive events was the same in this building, unlike Engineering Building-II where t_w of the successive events was double of the square events. The AEC of the successive test decreases compared to the square tests as seen in row 1 (right) of Figure 5. However, we notice an increase in ACC compared to the square events as seen in row 3 of Figure 5, indicating a more inefficient chiller response. We also notice a minimal impact on ECS for both square and successive events as seen in row 4 of Figure 5.

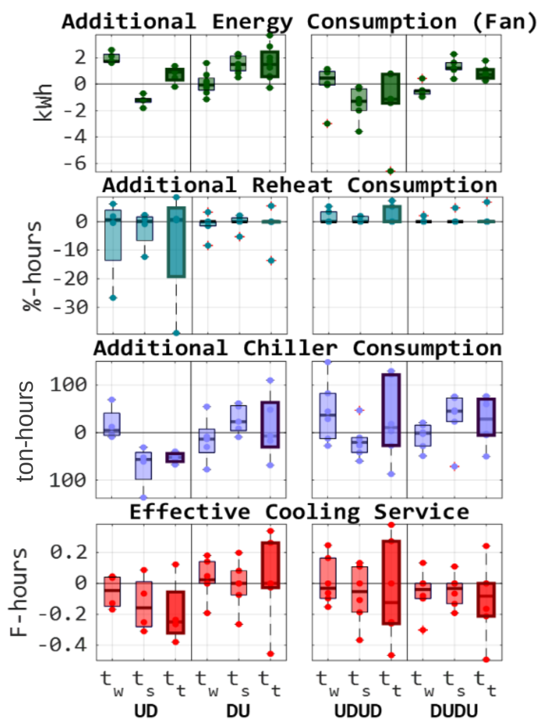

Figure 5: Efficiency metrics for fan, reheat, chiller and effective cooling service for square and successive signals for Thayer Building.

Summary of results

Table 2 presents a summary of the metrics computed for all eight buildings (we present the standard deviation associated with each metric in Section A.5 of the Electronic Appendix). For Engineering Building-III we see trends similar to Engineering Building-II. The Down-Up square tests have a higher mean AEC, and consequently a higher ACC value compared to the Up-Down square events. The AEC and ACC for the successive Up-Down-Up-Down test also increase compared to the Up-Down square and ramped test (the number of non-outliers for the Down-Up-Down-Up test was low). For Sullivan Shops-III we see an improvement in efficiency for the successive tests (lower AEC) compared to the square and ramped tests. Overall, the low AEC values indicate a minimal impact on the efficiency of the fan power response. For the Bob & Betty Beyster, Thayer, and Dana buildings we see a higher mean AEC for Down-Up events compared to Up-Down events, which is consistent with our prior experimental results. We observe very low AEC values for the North Quadrangle Complex indicating a minimal impact on the efficiency of fan power response.

Table 2. Summary table of metrics for all buildings. Values shown are for total window t_t

Building	**Shape: Polarity**[†]	**Fan Power: AEC, kWh**	**Chiller Power: ACC, gallons or ton-hours**	**Reheat Power: ARC, %-hours**	**Temp.Deviation: ECS, °F-hours**
Engineering Building-II	SQ: UD	-4.3	930	4.6	-0.11
	SQ: DU	11	3100	10	0.16
	RA: UD	0.17	-74	14	0.03
	RA: DU	8.8	440	-6.7	-0.09
	SU: UDUD	13	6000	0.84	-0.35
	SU: DUDU	35	4300	1.6	0.52
Engineering Building-III	SQ: UD	-0.66	900	-5.3	0.18
	SQ: DU	3.0	1200	1.1	0.28
	RA: UD	-1.8	1000	7.5	0.25
	RA: DU	1.1	2300	13	0.25
	SU: UDUD	3.6	3400	1.4	0.51
	SU: DUDU	6.8	790	2.7	0.86
Sullivan Shops- III	SQ: UD	0.18	N/A	12	0.42
	SQ: DU	0.17	N/A	6.2	0.71
	RA: UD	-0.22	N/A	4.9	0.33
	RA: DU	0.11	N/A	4.1	0.48
	SU: UDUD	-1.2	N/A	13	1.7
	SU: DUDU	-0.27	N/A	8.2	1.4
Park Shops	SQ: UD	-1.4	N/A	8.7	-0.22
	SQ: DU	-0.52	N/A	7.3	0.16
	RA: UD	N/A	N/A	1.0	0.00
	RA: DU	-0.20	N/A	11	-0.14
	SU: UDUD	-2.9	N/A	22	0.73
	SU: DUDU	0.52	N/A	15	-0.21
Bob & Betty Beyster Building	SQ: UD	-1.7	34	-9.8	-0.04
	SQ: DU	2.2	110	-22	0.11
	SU: UDUD	-7.2	-240	-22	0.01
	SU: DUDU	3.2	-94	-23	-0.11
Thayer Building	SQ: UD	0.72	-52	-7.3	-0.19
	SQ: DU	1.5	13	-1.4	0.02
	SU: UDUD	-1.5	26	2.1	-0.05
	SU: DUDU	0.84	25	1.2	-0.10
North Quadrangle Complex	SQ: UD	0.94	320	4.8	0.05
	SQ: DU	-0.52	-140	1.0	0.37
	SU: UDUD	-1.4	-160	5.2	-0.14
	SU: DUDU	-0.36	-210	3.7	0.31
Dana Building	SQ: UD	3.0	-16	4.7	-0.10
	SQ: DU	5.0	-16	-2.8	-0.30
	SU: UDUD	-1.3	N/A	N/A	N/A
	SU: DUDU	-0.44	N/A	N/A	N/A

[†] SQ: Square; RA: Ramped; SU: Successive; UD: Up-Down; DU: Down-Up; UDUD: Up-Down-Up-Down; DUDU: Down-Up-Down-Up

Conclusion

We experimented on eight buildings at NCSU and UM to quantify the holistic impacts on the different subcomponents of a commercial building HVAC system when it responds to load shifting grid services. In our previous work, we found that the vast thermal inertia of commercial building HVAC systems can be used to shift its power on minutely to hourly timescales to provide additional flexibility on the grid. We had also found that the polarity of setpoint signal to the building can have a significant impact on the efficiency of building response although the magnitude of impact varied from building to building. The key findings in this work are as follows:

- We found six out of eight of the buildings to respond consistently with our prior experimental results and find that the polarity of the setpoint signal has an impact on efficiency of fan power response.
- We observed a notable impact on terminal reheat systems when the buildings are used for minutely to hourly load shifting.
- Load shifting services can have a significant impact on the chiller power consumption.
- We find load shifting services to have a minimal impact on the effective cooling service provided to the thermal zones.
- Successive events cause a larger change in effective cooling service provided to the building. Changes to efficiency of fan power and chiller response varied from building to building and also depended on the testing window. More testing is needed to fully understand the impact of successive tests on efficiency of HVAC response.
- Ramping of setpoints can be used to improve the efficiency of building response although more experiments and analysis are needed to full quantify the impact.

This research is of significance to third party demand response aggregators for designing the optimal setpoint signals that will cause the building to closely follow the grid signal requested by the system operator. Our research offers additional insights into the additional building energy consumption that could be induced by participation, providing guidance on the necessary compensation for building operators. As the electrification of building heating systems gains traction (shifting away from natural gas), the interplay of HVAC subcomponent dynamics becomes even more important since the impacts on reheat can impact the building's ability to track the commanded power signal.

In Summer-Fall of 2020, we will continue experimenting on the buildings on both campuses. Our future experiments will be informed by a combination of the analysis of the experiments presented in this paper and modeling the experiments using building software like EnergyPlus and the Buildings Library in Modelica. This two-pronged approach will not only help us design optimal setpoint signals that mitigate inefficiency but also directly address the mismatch in the modeling and experimental work that has been observed in prior work.

Acknowledgments

This work was supported by the U.S. Department of Energy via the project IDREEM: Impact of Demand Response on short and long term building Energy Efficiency Metrics funded by the Building Technologies Office under contract number DE-AC02-76SF00515. We are immensely grateful to the facilities and operations staff of the UM and NCSU for supporting the research team in setting up the experiments. We thank Rishee Jain and Abigail Andrews for the insightful discussions. We also thank our Technical Advisory Group for providing their feedback to the research team.

Appendix

The electronic appendix for the paper can be found here: https://drive.google.com/file/d/1L16CYHr385IAjyEuXdENLTTkJzHCNDCl/view?usp=sharing

References

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.

Bender, J. 2019. “BACpypes”. *GitHub Repository*. [Online]. Link: https://github.com/JoelBender/bacpypes. Accessed on: 02/17/2020.

Brattle Group. 2019. “The National Potential for Load Flexibility”. [Online]. Link: https://brattlefiles.blob.core.windows.net/files/16639_national_potential_for_load_flexibility_-_final.pdf. Accessed on: 02/23/2020

Cai, J., J.E. Braun and Y. Lin. 2018. “Laboratory demonstration of variable-speed packaged air-conditioning units for provision of frequency regulation services.” *IEEE Power & Energy Society Innovative Smart Grid Technologies Conference*.

Callaway, D.S. and I.A. Hiskens. 2011. “Achieving controllability of electrical loads,” Proceedings of the IEEE, 99(1): 184-199.

Davis, S.J., N.S. Lewis, M. Shaner, S. Aggarwal, D. Arent, I.L. Azevedo, S.M. Benson, et al. 2018. “Net-zero emissions energy systems.” *Science*, *360*(6396), p.eaas9793.

Fan, C. and Y. Ding. 2019. “Cooling load prediction and optimal operation of HVAC systems using a multiple nonlinear regression model.” *Energy and Buildings*, 197: 7-17.

Goddard, G., J. Klose, and S. Backhaus. 2014. “Model development and identification for fast demand response in commercial HVAC systems.” *IEEE Transactions on Smart Grid*, 5(4): 2084–2092.

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): 2066-2014.

Keskar, A., D. Anderson, J.X. Johnson, I.A. Hiskens and J.L. Mathieu. 2018. "Experimental investigation of the additional energy consumed by building HVAC systems providing grid ancillary services." *ACEEE Summer Study on Energy Efficiency in Buildings.*

Keskar, A., D. Anderson, J.X. Johnson, I.A. Hiskens and J.L. Mathieu. 2019. "Do commercial buildings become less efficient when they provide grid ancillary services?" *Energy Efficiency* https://doi.org/10.1007/s12053-019-09787-x

Lei, S., J.L. Mathieu, and R.K. Jain. 2019. "Performance of existing baseline models in quantifying the effects of short-term load shifting of campus buildings." *SLAC-R-1131*.

Lin, Y., P. Barooah, S. Meyn, and T. Middelkoop. 2015. "Experimental evaluation of frequency regulation from commercial building HVAC systems." *IEEE Transactions on Smart Grid*, 6(2): 776-783.

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.

Mathieu, J.L., A.J. Gadgil, D.S. Callaway, P.N. Price and S. Kiliccote. 2010. "Characterizing the response of commercial and industrial facilities to dynamic pricing signals from the utility." *In: Proceedings of ASME 4th International Conference on Energy Sustainability*.

Raman, N.S. and P. Barooah. 2018a. "On the round-trip efficiency of an HVAC-based Virtual Battery". https://arxiv.org/abs/1803.02883.

Raman, N.S., and P. Barooah. 2018b. "Analysis of Round-Trip Efficiency of an HVAC-Based Virtual Battery". *International High Performance Buildings Conference*. Paper 287. Available: https://docs.lib.purdue.edu/ihpbc/287

Watson, D.S, S. Kiliccote, N. Motegi, and M.A. Piette. 2006. "Strategies for Demand Response in Commercial Buildings." *ACEEE Summer Study on Energy Efficiency in Buildings*.

Zhao, P., G.P. Henze, S. Plamp, V.J. Cushing. 2013. "Evaluation of commercial building HVAC systems as frequency regulation providers." *Energy and Buildings*. 67: 225–235.