Experimental investigation of the additional energy consumed by building HVAC systems providing grid ancillary services

Aditya Keskar, University of Michigan
David Anderson, University of Michigan
Jeremiah X. Johnson, North Carolina State University
Ian A. Hiskens, University of Michigan
Johanna L. Mathieu, University of Michigan

ABSTRACT

Commercial buildings offer vast thermal energy storage capability. Control of building heating, ventilation, and air conditioning (HVAC) systems can potentially be used to balance variations in renewable generation and load. Specifically, buildings can provide ancillary services to the grid by decreasing and increasing consumption with respect to their baseline, making them appear as a battery to the power system operator. However, a recent study has shown that buildings providing these services tend to consume more energy, resulting in a low effective round-trip efficiency. To explore this phenomenon further, experiments were conducted on three buildings on the University of Michigan campus. The buildings were chosen to represent a variety in structure, size, and HVAC system layout. They were instrumented in early summer of 2017 and baseline power and building automation system (BAS) data were collected for several months. The building thermostats were then perturbed through predefined patterns emulating ancillary service events, enabling detailed investigation of the resulting electrical energy consumption. This paper presents experimental results, focusing on the additional energy consumed and effective input/output efficiency. We find that the efficiency of building response depends on the magnitude and polarity of the temperature setpoint changes. Our results are consistent with past experimental results, but inconsistent with past modelling results. This indicates that the models need to be improved in order to capture the energy impacts of ancillary service provision by buildings.

Introduction

The increasing share of renewable energy on the grid presents greater reliability challenges to the power system operator. Responding to the stochastic nature of wind and solar power production will require new, creative, and efficient solutions. Traditionally, ancillary services have been provided by conventional generation resources that ramp up or down to maintain the crucial supply-demand balance required to maintain grid frequency. In recent years, there has been a significant amount of research done in providing these services using electrical loads (Callaway and Hiskens 2015).

Commercial building HVAC (heating, ventilation, and air conditioning) systems present an immense energy-storage resource that could be harnessed to maintain supply-demand balance on the grid. Commercial buildings account for roughly 20% of the energy consumed in the United States (US Energy Information Administration 2016). Buildings can provide ancillary services by increasing and decreasing building power consumption with respect to their baselines in response to signals sent from the power system operator (Beil et al. 2015).

A recent experimental study conducted at Los Alamos National Laboratory (LANL) (Beil et al. 2015) showed that when buildings are subjected to power perturbations emulating ancillary services, they tend to consume additional energy over their baseline even when the ancillary service events are designed to achieve only load shifting (i.e., no net change in energy consumption over the event). Specifically, the ancillary service tests designed by Beil et al. (2015) first increase/decrease building temperature setpoints (decreasing/increasing power consumption) for a fixed length of time and then decrease/increase them (increasing/decreasing power consumption) by the same amount for the same length of time. The study found that the average effective round-trip efficiency (RTE) of the building was only 46%, which is significantly lower than standard storage technologies like redox flow batteries (75%), lithium-ion batteries (80%, or higher), and pumped hydro storage (81%) (Pacific Northwest National Laboratory 2013).

Using a physics-based model, the LANL experiments were simulated in Lin et al. (2017) in order to understand the cause for the inefficiency. However, the simulation results were not universally consistent with the experimental results from Beil et al. (2015). While the simulation results predicted an RTE of less than 1 when power is increased above the baseline and then subsequently decreased (we call this an Up-Down power variation), it predicted an RTE of greater than 1 when power is decreased first below its baseline then increased (a Down-Up power variation). The change in energy consumption was attributable to a change in average room temperatures. In contrast, the experimental results always showed an RTE of less than 1 and showed that the RTE associated with Up-Down power variation is *higher* than the RTE associated with Down-Up power variation. Further experimental results were deemed necessary to explore this phenomenon.

This paper presents results from a series of similar experiments conducted on three campus buildings at the University of Michigan. We first describe the experiment setup with a detailed description of the campus buildings and metrics used to quantify the efficiency of building response to ancillary service tests. We then present the results of the experiments and a discussion of the response of each building's HVAC system, focusing on the difference in building response to an Up-Down power variation versus a Down-Up power variation. We conclude by emphasizing the need to gather further experimental evidence and develop better models that accurately capture the response of the buildings.

Experimental Setup

Figure 1 shows the ancillary services events, which are identical to those in Beil et al. (2015). During the test window (t_w), temperature setpoints are increased and decreased symmetrically, where full magnitude (FM) refers to the change in temperature setpoint at the midpoint of the event. Increases/decreases in temperature setpoints decrease/increase fan power consumption. Specifically, when room temperature setpoints are decreased in a room with a controller in cooling mode, the VAV (variable air volume) damper opens to increase air flow to the room. This causes a drop in static pressure, and the control system responds by increasing the speed of the fan associated with that room's zone, thereby increasing the fan's power consumption. Similarly, when room temperature setpoints are increased, the VAV damper modulates to decrease the air flow, resulting in reduced fan power consumption.

We refer to the ordering of the power increase and decrease as the polarity of the event. Due to a building's large thermal inertia, room temperatures do not change significantly during an ancillary service event. Note that previous studies (Hao et al. 2014) have used direct fan power control to vary the power consumption of buildings. In contrast, we use temperature setpoint control since temperature setpoints can be readily manipulated within Building

Automation Systems (BAS). Temperature setpoint adjustment allows us to work with the BAS whereas direct manipulation of fan power consumption acts as a disturbance to the BAS. The BAS will seek to counteract that disturbance by opening/closing VAV dampers, which eventually restores the fan power consumption to its previous value.

The change in energy consumption with respect to the baseline is estimated for the test window and settling window (t_s), which comprise the total window ($t_t = t_w + t_s$). We use a test window t_w = 1 hour. Beil et al. (2015) used t_w = 30 minutes; however, one of the buildings we use in our experiments (Weill Hall) did not exhibit clear responses when we attempted t_w = 30 minutes, so we used t_w = 1 hour for consistency across campus buildings.

Beil et al. (2015) defined the RTE as the ratio of the energy consumed by the building below its baseline over t_t to the energy consumed by the building above its baseline over t_t (see Fig. 3). They found the mean RTE of an Up-Down power variation to be 61% and mean RTE of a Down-Up power test to be 34%. Conversely, the simulation work by Lin et al. (2017) found the mean RTE of an Up-Down power variation to be 88% and a Down-Up power variation to be 109%.

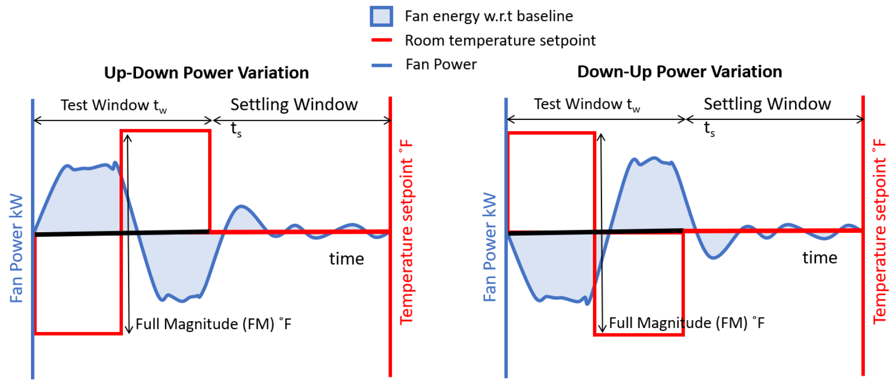

Figure 1. Up-Down (left) and Down-Up (right) power variation resulting from temperature setpoint changes.

The Buildings

We conducted experiments on three buildings (shown in Fig. 2 along with their average load profiles) on the University of Michigan campus. The buildings vary in terms of size, structure, and HVAC system layout as shown in Table 1.

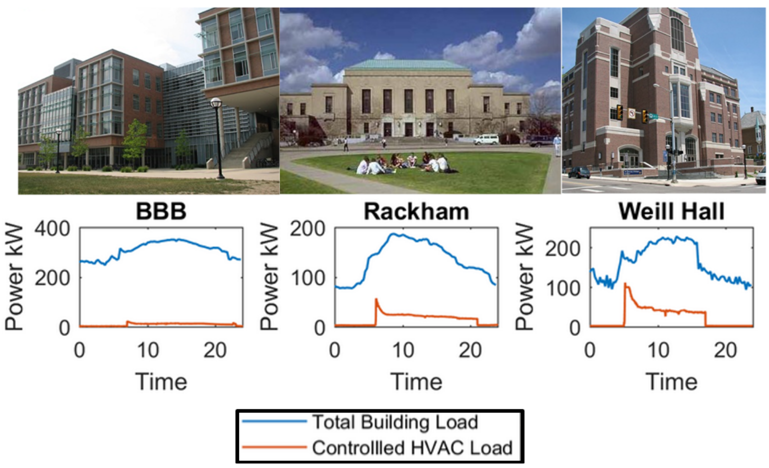

Figure 2. Bob & Betty Beyster Building (BBB) (left), Rackham Building (center), and Weill Hall (right).

Table 1. Building parameters and testing information.

	Rackham	**BBB**	**Weill Hall**
Year of construction	1938	2005	2006
Type	Office/Auditorium	Classroom/Office	Classroom/Office
Area	157,957 ft^2	104,132 ft^2	97,989 ft^2
2016 Energy Consumption	972 MWh	3160 MWh	1030 MWh
2016 Peak Demand	226 kW	391 kW	352 kW
Electricity provider	Campus combined heat & power plant	DTE Energy (electric utility)	Campus combined heat & power plant
HVAC	8 AHUs, 8 supply fans, 8 return fans	3 AHUs, 4 supply fans, 3 return fans	2 AHUs, 2 supply fans, 2 return fans
# of AHU zones controlled	4	1	2
# of setpoints controlled	109	193	104
# of fans instrumented	8 (4 supply fans and 4 return fans)	2 (1 supply fan and 1 return fan)	5 (2 supply fans, 2 return fans, 1 cooling tower fan)
Baseline months (Unperturbed HVAC operation)	August and October 2017	July, August, and October 2017	June, July, August, and October 2017
Test months	September 2017	September 2017	September 2017
Test hours (2 tests per day)	9:00-10:00 1:00-2:00 (weekdays)	18:00-19:00 21:00-22:00 (everyday)	9:00-10:00 1:00-2:00 (weekdays)

Measuring Fan Power Consumption

Each building has multiple air handling units (AHUs), each serving different zones of the building. In each building, we controlled the temperature setpoints associated with a subset of the zones. We installed current sensors on a single phase of the power lines serving the fans associated with the controlled zones. We assume constant power factors and voltages (determined using one week of measured voltage and power factor data) and use these values to compute the three-phase fan power. Table 2 shows the power and air flow ratings of the fans whose power was monitored.

Table 2. Ratings of the fans that were monitored in each building.

Building	AHU	Power (HP/kW)		Air flow (1000 cfm)	
		Return Fan	Supply Fan	Return Fan	Supply Fan
Rackham	2	20/14.9	10/7.5	23.6	20
	4	20/14.9	10/7.5	23	20
	7	30/22.4	15/11.2	24	23
	8	30/22.4	15/11.2	24	23
BBB	1	60/44.7	20/14.9	36	32.6
Weill Hall (AHU 1 & 2)	1	100/74.6	40/29.8	45	43.5
	2	100/74.6	40/29.8	45	41.5

We also received minute-resolution data from the BAS of the air flow associated with each supply fan and the temperature of one room per AHU zone. Room temperatures were used to gauge potential occupational discomfort caused by the tests.

Efficiency Metrics

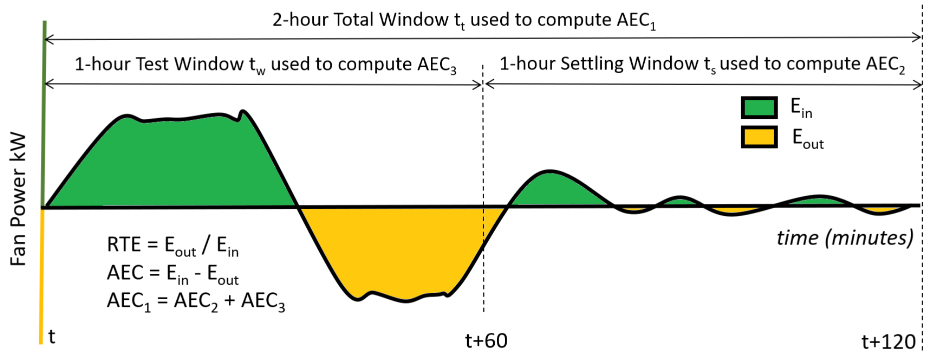

Figure 3. Metrics used to quantify the building response efficiency, where E_{in} and E_{out} are both defined as positive.

Figure 3 shows the metrics we use to quantify building response efficiency. In addition to the RTE, we also define the Additional Energy Consumption (AEC), which is the energy consumed by the building above its baseline (E_{in}) minus the energy consumed by the building below its baseline (E_{out}). For a particular building, a higher AEC indicates lower RTE. The AEC during the full window t_t is denoted by AEC_1, while the AEC during the test window (t_w) and during the settling window (t_s) is denoted by AEC_3 and AEC_2, respectively. Since we design the tests to achieve only load shifting over t_w, we would hope $AEC_3 = 0$. However, the tests are imperfect, so the metric generally takes a non-zero value, varying from event to event. As shown in the figure, our goal was to use a settling window of t_s = 1 hour, but due to a daily dip in fan power consumption at 11:00 A.M. at Weill Hall and the HVAC switching to night-time operation at 11:00 P.M. in the BBB building, the settling window t_s for both buildings was reduced to 48 minutes.

Outliers

We filter out tests that are not sufficiently responsive to temperature setpoint changes. A potential reason for these outliers could be low occupancy, which can cause the fan power to be unaffected by setpoint changes. Another possible reason is a lack of response from VAV boxes that are already operating at their maximum or minimum air flow capacity during the testing window. These boxes cannot respond any further to achieve the commanded setpoint. We deem a test an outlier if it does not satisfy *both* of the following two criteria:

1. During the testing window t_w, the response should be sufficiently symmetric. Specifically, if $E_{out} > E_{in}$, then E_{in} should be at least 20% of E_{out} and if $E_{in} > E_{out}$, then E_{out} should be at least 20% of E_{in}.

2. During the testing window t_w, the response should be sufficiently large. Specifically, $E_{out}+E_{in}$ should be above a tolerance, which is tuned separately for each building. The tuning process is heuristic: we tune the tolerances based on visual inspections of the time series data. We use the following tolerances: Weill Hall: 4.5 kWh, Rackham Building: 3 kWh, and BBB:

1.5 kWh. We investigate the sensitivity of the efficiency metrics to the tolerances in the results section.

Baseline Estimation

Beil et al. (2015) estimated the baseline by linearly interpolation, with the ends of the linear baseline estimate given by fan power data over short time windows just prior to and after the ancillary service event. Afshari et al. (2017), used air flow and room temperature data to develop algorithms that compute when the building settles back to baseline operation after an ancillary service event.

We follow a similar methodology to Beil et al. (2015). Specifically, we use least squares to fit a linear baseline to the fan power data over the 5 minute period just before the event and the 5 minute period immediately after the settling window. Subtracting the baseline from the HVAC power consumption gives the estimated change in power consumption over the full window t_t.

Baseline estimates often have significant error (Mathieu et al. 2011). To evaluate the baseline error, we test the baseline method on days in which there were no ancillary service events (referred to as baseline days). Specifically, we use our baseline method to compute AEC_1 at the building-specific event times on baseline days in August and October. Since no actual tests were conducted on baseline days, an accurate baseline method would result in $AEC_1 = 0$. Figure 4 shows boxplots of AEC_1 for each building. We use these results to compute 95% confidence intervals on mean AEC_1 estimates. Since we filter ancillary service event outliers, we also filter baseline outliers. Specifically, we remove baseline day AEC_1 estimates below the 5^{th} percentile and above the 95^{th} percentile. We use the remaining data to calculate the bias and standard deviation σ. Because the mean AEC_1 is calculated from n events, the corresponding standard deviation for mean AEC_1 is $\sigma/\sqrt{n}$. We therefore define the 95% confidence interval (CI) for the mean AEC_1 as $\pm 1.96\ \sigma/\sqrt{n}$. The bias and CI for each mean AEC_1 are presented together with the associated results in the following tables.

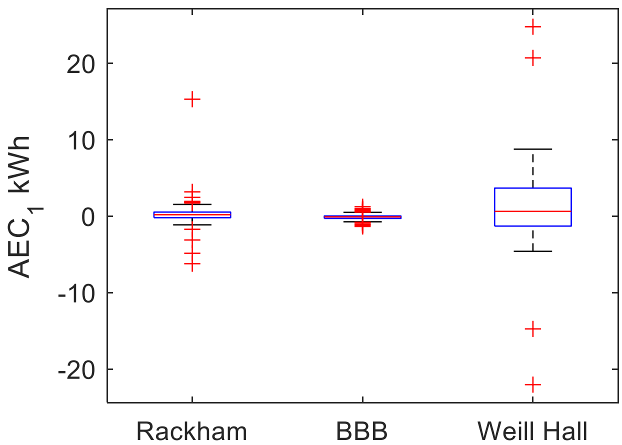

Figure 4. Distribution of AEC_1 for baseline days in August and October, for each building. The red line within the box shows the median of AEC_1 values. The bottom and top edges of the box mark the 25^{th} and 75^{th} percentile, respectively. The whiskers extend to the most extreme data points deemed non-outliers. The outliers are shown with red pluses.

Results and Discussion

Bob & Betty Beyster (BBB) Building

A total of 52 tests were conducted on the BBB Building. Figure 5 shows HVAC power consumption and room temperature for three representative events. As seen in the figure, there is no significant change in room temperature. Table 3 shows the results for four event types (varying in polarity and FM). We conducted 13 tests of each type. The table shows the number of non-outliers n used to compute the efficiency metrics. If n < 5, we do not present the standard deviation of the metrics.

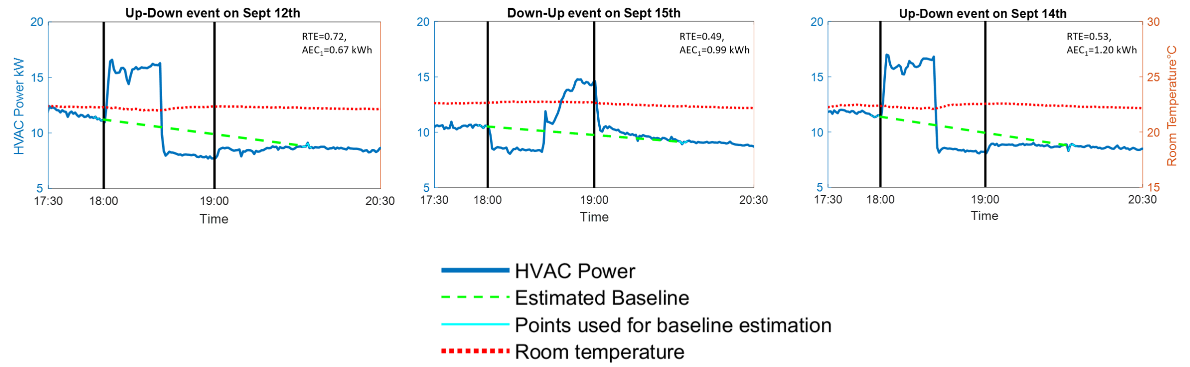

Figure 5. Three representative events at the BBB Building including HVAC power consumption (left axis) and room temperatures (right axis).

Table 3. Efficiency Metrics for the Bob & Betty Beyster Building

Event Type		1	2	3	4
Power Polarity		Up-Down	Up-Down	Down-Up	Down-Up
Full Magnitude (FM)		2˚F	4˚F	2˚F	4˚F
n, # of non-outliers (outliers)		3 (10)	9 (4)	3 (10)	7 (6)
RTE: Mean (σ)		1.16 (N/A)	0.79 (0.20)	2.46 (N/A)	0.47 (0.12)
AEC_1 (kWh)	Mean (σ)	-0.098 (N/A)	0.47 (0.47)	-0.86 (N/A)	0.99 (0.45)
	Bias ±95% CI	-0.13±0.35	-0.13±0.20	-0.13±0.35	-0.13±0.22
AEC_2 (kWh): Mean (σ)		-0.34 (N/A)	-0.58 (0.07)	-0.11 (N/A)	0.50 (0.16)
AEC_3 (kWh): Mean (σ)		0.24 (N/A)	1.05 (0.44)	-0.75 (N/A)	0.50 (0.31)

First, we compare Event Types 2 and 4, for which FM = 4 ˚F. Up-Down power variation (Event Type 2) averaged a higher RTE (79%) than Down-Up power variation (Event Type 4), which averaged 47%. This is consistent with the results from Beil et al. (2015). The increased efficiency of Up-Down variation over Down-Up variation is further validated by the lower mean AEC_1 of 0.47 kWh (Up-Down) compared with 0.99 kWh (Down-Up). AEC_3 and AEC_2 give us further insight into the building's response. Up-Down power variation consumes more energy in the test window (see AEC_3) but less energy in the settling window (see AEC_2) than Down-Up power variation. Thus, overshoot after the test makes the Down-Up power variation more inefficient than Up-Down power variation.

Event Types 1 and 3, for which FM = 2 ˚F, have mean RTEs of 116% (Up-Down) and 246% (Down-Up), along with negative mean AEC_1 values. These results indicate a *decrease* in building energy consumption resulting from the provision of ancillary services. However, for both of these event types we have removed a substantial number of outliers since the building often appeared nonresponsive to FM = 2˚F tests. Therefore, more testing is needed to verify these efficiency metric values.

Weill Hall

Eighteen tests were conducted on Weill Hall. All tests used FM = 2 ˚F as there was a concern that the occupants would notice FM = 4 ˚F tests. Figure 6 shows HVAC power consumption and room temperature for three representative events. Weill Hall was the only building with occupational sensors installed which impacted the building's response to the tests.

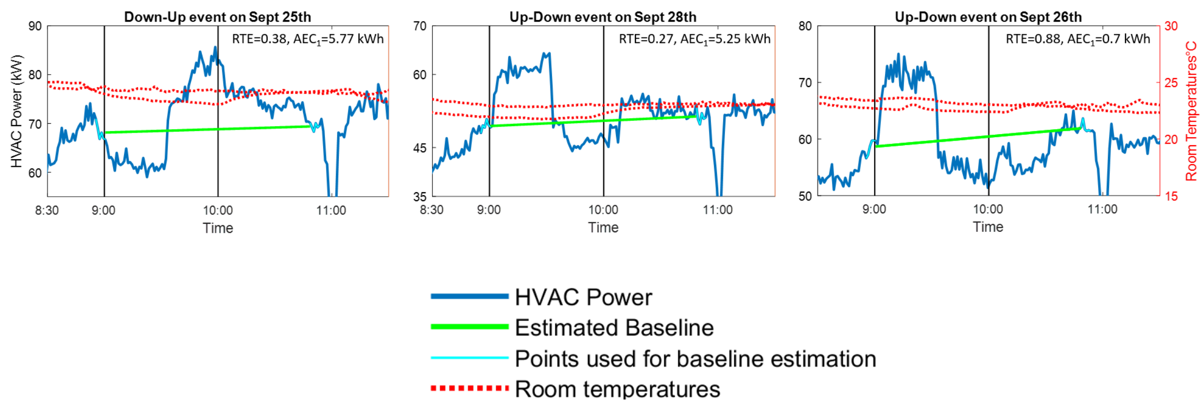

Figure 6. Three representative events at Weill Hall including HVAC power consumption (left axis) and room temperatures (right axis).

Table 4 shows the results of two event types (a total of 9 events of each type were conducted). The mean RTE of Up-Down power variation (68%) was higher than the Down-Up power variation (34%). The AEC_1 metric also showed that the Up-Down power variation consumed less energy on average (1.98 kWh) than the Down-Up power variation (5.07 kWh) over its baseline. Again, Up-Down power variation consumes more energy in the test window (see AEC_3) but less in the settling window (see AEC_2) than Down-Up power variation.

Table 4. Efficiency Metrics for Weill Hall

Event Type		1	2
Power Polarity		Up-Down	Down-Up
Full Magnitude (FM)		2˚F	2˚F
n, # of non-outliers (outliers)		6 (3)	6 (3)
RTE	Mean	0.68 (0.43)	0.34 (0.16)
AEC_1 (kWh)	Mean (σ)	1.98 (2.28)	5.07 (2.28)
	Bias ±95% CI	1.1±2.28	1.1±2.28
AEC_2 (kWh): Mean (σ)		0.04 (1.90)	3.33 (1.48)
AEC_3 (kWh): Mean (σ)		1.95 (2.26)	1.79 (1.60)

Rackham Building

A total of 32 tests were conducted on the Rackham Building. Figure 7 shows HVAC power consumption and room temperature for three representative events. Table 5 gives the efficiency metrics by event type. Overall, we conducted 9 events (of full magnitude 2˚F) and 7 events (of full magnitude 4˚F) for each polarity in the building.

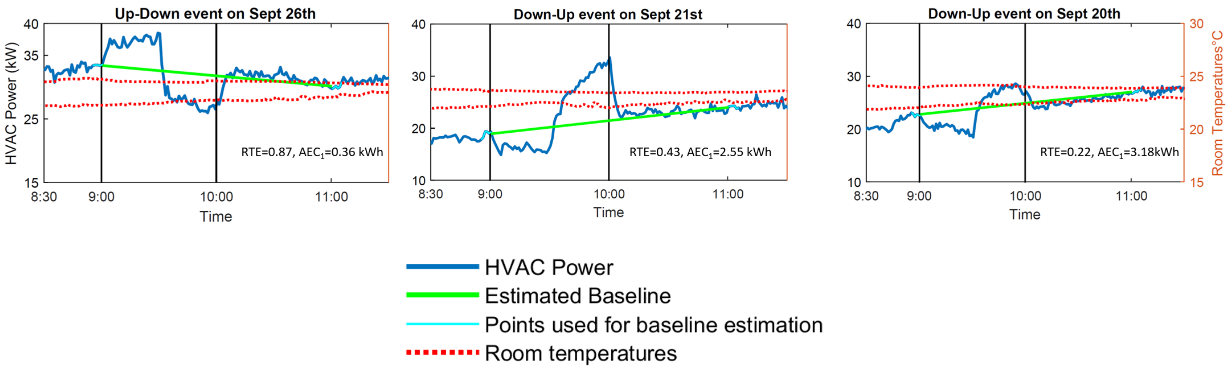

Figure 7. Three representative events at the Rackham Building including HVAC power consumption (left axis) and room temperatures (right axis).

Table 5. Efficiency metrics for the Rackham Building

Event Type		1	2	3	4
Power Polarity		Up-Down	Up-Down	Down-Up	Down-Up
Full Magnitude (FM)		2˚F	4˚F	2˚F	4˚F
n, # of non-outliers (outliers)		2 (7)	4 (3)	2 (7)	5 (2)
RTE	Mean (σ)	0.99 (N/A)	0.81 (N/A)	1.21 (N/A)	0.49 (0.41)
AEC_1 (kWh)	Mean (σ)	0.85 (N/A)	0.77 (N/A)	-0.30 (N/A)	2.12 (1.60)
	Bias ± 95% CI	0.21±0.89	0.21±0.63	0.21±0.89	0.21±0.56
AEC_2 (kWh): Mean (σ)		0.33 (N/A)	-0.91 (N/A)	-0.16 (N/A)	1.11 (0.51)
AEC_3 (kWh): Mean (σ)		0.53 (N/A)	1.67 (N/A)	-0.14 (N/A)	1.01 (1.23)

Comparing Event Types 2 and 4, Up-Down power variation averaged an RTE of 81% and Down-Up power variation averaged an RTE of 49%. The mean AEC_1, AEC_2, and AEC_3 values follow the same trends as seen in the other buildings. Similar to the BBB Building, we have removed a substantial number of outliers corresponding to Event Types 1 and 3 since the building often appeared nonresponsive to FM = 2˚F tests. Again, more testing is needed to verify these efficiency metric values.

How do the buildings compare?

Figures 8 and 9 summarize the efficiency metrics across the three buildings (though only the FM = 4˚F tests are shown for BBB and Rackham). Figure 8 shows boxplots of the RTEs along with the mean RTEs for Up-Down and Down-Up events for each building, and the percentage decrease in mean RTE. Similarly, Figure 9 shows boxplots of the AEC_1 along with the mean AEC_1, and the percentage increase in mean AEC_1. In all buildings the mean RTE of Up-Down power variation is higher than the mean RTE of Down-Up power variation (Figure 8), and the opposite is true for the mean AEC_1 (Figure 9). Hence, the polarity of the ancillary service event plays a significant role in the building response efficiency. This is consistent with the results found in Beil et al. (2015).

We believe a potential reason that the Down-Up power variation is more inefficient is that the buildings over-respond to the setpoint changes in the middle of the test window. As the temperature setpoints are bought down to increase fan power consumption, the fans sharply increase their power consumption since the rooms have become warmer due to an increase in temperature setpoints during the previous half of the test.

Figure 10 shows the sensitivity of the mean RTE and AEC_1 in each building and each polarity to the outlier tolerances. Specifically, we calculate the metrics using 1) all tests, 2) only tests that satisfy the two conditions given earlier and using the default tolerances, 3) only tests that satisfy the two conditions and using a 10% increase in tolerance, and 4) only tests that satisfy the two conditions and using a 10% decrease in tolerance.

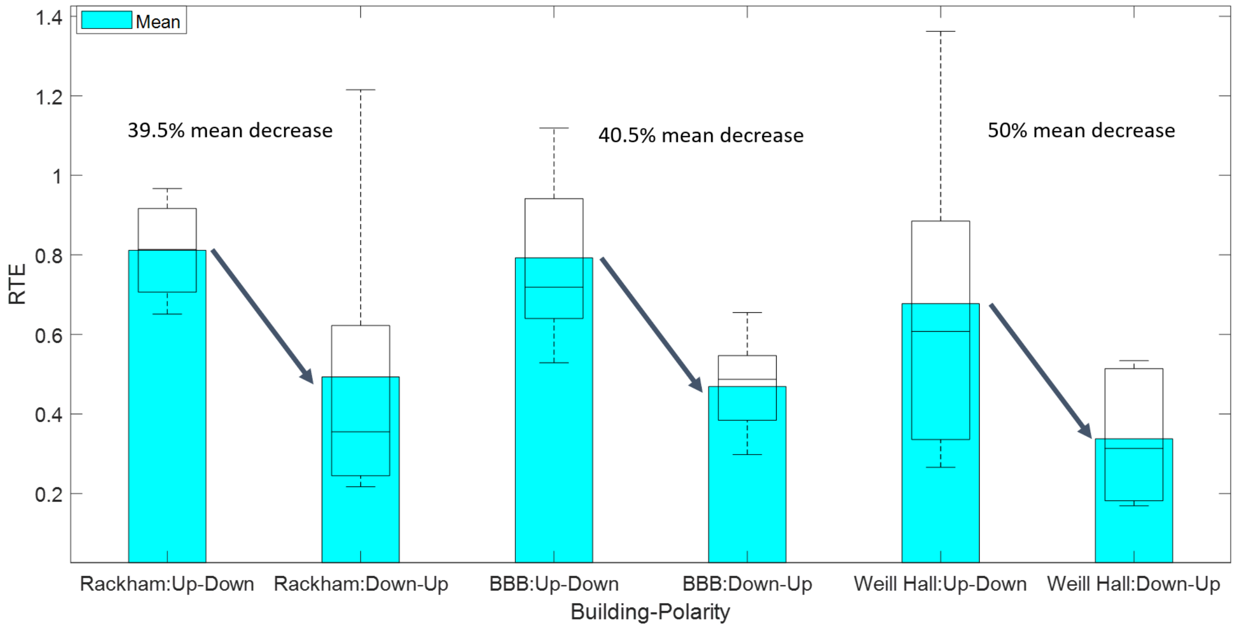

Figure 8. RTE statistics for all buildings.

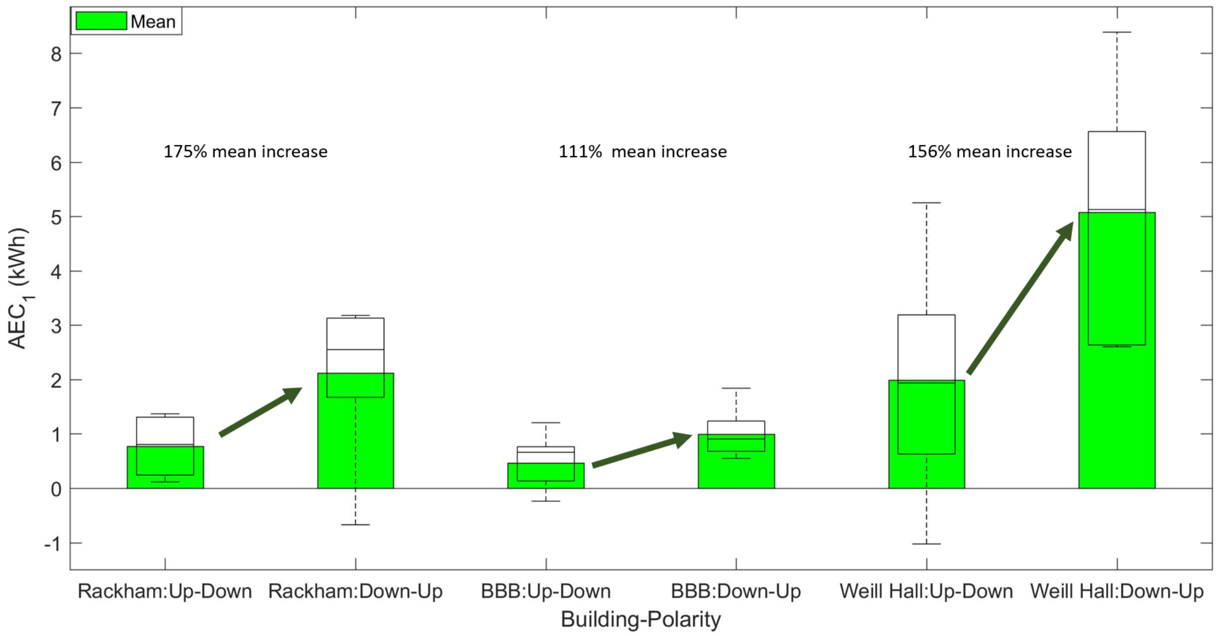

Figure 9. AEC_1 statistics for all buildings.

We find that while the tolerances do affect the results, the trends remain the same with the exception of Rackham. At Rackham, without removing outliers (red line), the Down-Up power variation is more efficient that the Up-Down power variation. This trend reversal is due to a single spurious test. When inspecting the time series data associated with this event we found that the building did not clearly respond to the setpoint changes and E_{in} was very small, giving a "divide by small number" problem.

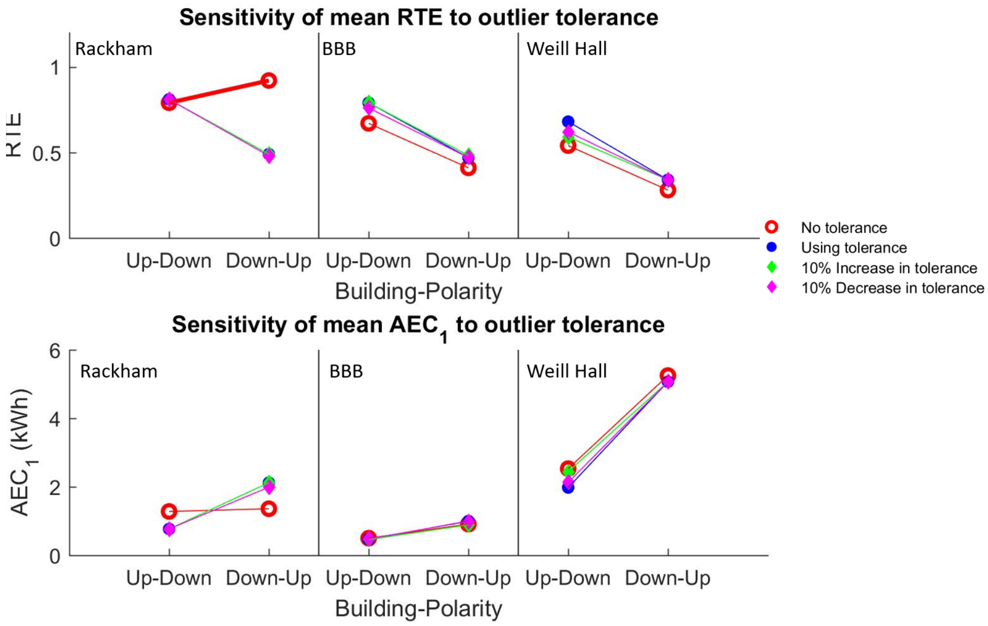

Figure 10. Evaluating sensitivity of results to change in tolerance.

Conclusion and Future Work

We conducted a total of 122 tests on three campus buildings at the University of Michigan in order to quantify the efficiency of building responses to ancillary service events. Our key findings are as follows.

- The vast thermal inertia of buildings can be used to provide ancillary services by varying HVAC power consumption on a short time scale with minimal occupant discomfort. Larger temperature setpoint changes cause a larger variation in fan power consumption.
- Up-Down power variations are more efficient than Down-Up power variations, which is consistent with the result of Beil et al. (2015). Thus, the polarity of the ancillary service event has a significant impact on the building response efficiency.
- The experimental results obtained in this paper and in Beil et al. (2015) are inconsistent with the modelling results of Lin et al. (2017). Better models are needed to capture the energy impacts of ancillary service provision by buildings.
- The AEC is a more robust metric than the RTE since it does not have "divide by small number" problems. However, it is hard to compare AECs across buildings.

The factors driving the inefficiency of building responses are still largely unknown. Our future work will develop better models that capture the building behaviour that we see in the experiments. Furthermore, we would like to conduct additional experiments, for example, solely Up-Power variation and solely Down-Power variation, which would help isolate the behaviour associated with each of the two phases of the symmetric Up-Down and Down-Up tests. We would also like to explore successive ancillary service tests, as explored in a recent simulation study (Raman and Barooah, 2017) that used a physics-based model to show the convergence of the RTE to unity as a building is subjected to more and more successive tests.

Acknowledgments

This research was funded by the University of Michigan's MCubed Program and a Rackham Summer Award. We are immensely grateful to Dr. Sina Afshari for setting up the experiments on campus. We also thank Paul Giessner for helping setup the code to process the building data, Scott Backhaus for insightful discussions, and Jonas Kersulis for his inputs on the paper. This work would not have been possible without the tremendous support of the staff at the University of Michigan's Facilities and Operations Group and Energy Management Group. We would especially like to thank Kevin Morgan, Andrew Cieslinski, and Connor Flynn from the Energy Management group for their constant support and feedback. We also thank the building managers of Rackham, BBB and Weill Hall for their cooperation with the research.

References

Afshari, S., J. Wolfe, M.S. Nazir, I.A. Hiskens, J.X. Johnson, J.L. Mathieu, Y. Line, A.K. Barnes, D.A. Geller and S.N Backhaus. 2017. "An Experimental Study of Energy Consumption in Buildings Providing Ancillary Services". *IEEE Power and Energy Society Innovative Smart Grid Technologies Conference.*

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.

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

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.

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, D.S. Callaway, and S. Kiliccote. 2011. "Variability in automated responses of commercial buildings and industrial facilities to dynamic electricity prices." *Energy and Buildings,* 43(12): 3322-3330.

Pacific Northwest National Laboratory (PNNL). 2013. "National Assessment of Energy Storage for Grid Balancing and Arbitrage. Volume 2: Cost and Performance Characterization." PNNL-21388: Phase II/Vol.2.

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

U.S Energy Information Administration (EIA). 2016. "Use of Energy in the United States Explained." https://www.eia.gov/energyexplained/.