

Lawrence Berkeley National Laboratory

LBL Publications

Title

A Critical Exploration of the Efficiency Impacts of Demand Response From HVAC in Commercial Buildings

Permalink

<https://escholarship.org/uc/item/4rp474bw>

Journal

Proceedings of the IEEE, 108(9)

ISSN

0018-9219

Authors

MacDonald, Jason S
Vrettos, Evangelos
Callaway, Duncan S

Publication Date

2020-09-01

DOI

10.1109/jproc.2020.3006804

Peer reviewed

A Critical Exploration of the Efficiency Impacts of Demand Response from HVAC in Commercial Buildings

Jason S. MacDonald
*Energy Storage and Distributed Resources
Lawrence Berkeley National Laboratory
Berkeley, CA, USA
jsmacdonald@lbl.gov*

Evangelos Vrettos
*Research and Digitilisation
Swissgrid
Aarau, Switzerland
evangelos.vrettos@swissgrid.ch*

Duncan S. Callaway
*Energy and Resources Group
University of California, Berkeley
Berkeley, CA, USA
dcal@berkeley.edu*

Abstract—Increasing quantities of renewable energy generation has yielded a need for greater energy storage capacity in power systems. Thermal storage in variable air volume (VAV) heating, ventilation, and air conditioning (HVAC) in commercial buildings has been identified as a possibly inexpensive source of grid storage, but the true costs are not known. Recent literature explores the inefficiency associated with providing grid services from these HVAC-based demand response resources by employing a battery analogy to calculate round-trip efficiency (RTE). Results vary significantly across studies, and in some cases reported efficiencies are strikingly low. This paper has three objectives to address these prior results. First, we synthesize and expand on insights in existing literature through systematically exploring the potential causes for the discrepancies in results. We reinforce previous work indicating baseline modeling may drive differences across studies, and deduce that control accuracy plays a role in the major differences between experiments and simulation. Second, we discuss why the RTE metric is problematic for demand response applications, discuss another proposed metric, additional energy consumption (AEC), and propose an extension, which we call uninstructed energy consumption (UEC), to evaluate demand response performance. Finally, we explore the merits of different metrics using experimental data and highlight UEC’s reduced sensitivity to the characteristics of the demand response signal than previously proposed metrics.

Index Terms—Demand Response, Battery Analogy, Round-trip Efficiency, Energy Cost

I. INTRODUCTION

The increased deployment of variable, supply-limited renewable generation resources into the power grid has highlighted a need for energy storage and demand side resources to support balancing operations in power systems [1]. Energy storage – including batteries and pumped hydropower – is capable of absorbing renewable generation intermittency, improving power quality, supporting competitive markets through price arbitrage, and generally providing highly flexible capacity to manage a myriad of system needs [2], [3]. However, the

high capital cost and siting restrictions for most energy storage technologies has had regulators, industry, and the research community looking for alternative resources to manage the challenges of the future generating resource mix.

Demand response (DR) is one such alternative storage resource of interest. DR has a history as emergency capacity and as spinning reserve in many balancing areas in the US [4], [5]. In wholesale and retail electricity markets, DR participation is growing [6]. State [7] and federal [8] regulators seek opportunities to open markets to DR and adopt rules to support the removal of participation barriers. While a considerable body of work exists discussing DR control techniques to provide grid services [9]–[13], identification of the costs of participation for controllable load in DR programs has been a challenge, with the vast majority of the work focused on evaluating response to time-of-use pricing [14].

Recent work has focused on incremental energy costs in buildings providing DR, with a focus on heating, ventilation and air conditioning (HVAC) systems in commercial buildings. Some papers have identified the additional energy consumption caused by a system diverging from an energy optimal demand schedule to hold reserve capacity [15]–[17], while others focus on the additional energy consumption caused by responding to grid control signals [15], [16], [18]–[23]. This latter group of work employs an analogy between thermal storage in HVAC systems and chemical storage in batteries to evaluate round-trip efficiency (RTE), or the ratio of discharging to charging energy. In [18] an illustrative square wave DR event was used to calculate the round trip efficiency, and versions of this service were reproduced in a number of follow-on studies. These studies – which include both experiments and simulations – report widely varying results, from below 50% RTE [18], [20] to results near 100% or better [16], [19], [21]. The apparent lack of agreement between these works on the efficiency of demand response makes interpretation difficult.

This paper explores the methodological similarities and differences across these conflicting works to support a convergence of understanding. We build upon their insights and

The work described in this study was coordinated by the Consortium for Electric Reliability Technology Solutions (CERTS), and was funded by the U.S. Department of Energy’s Office of Electricity Delivery and Energy Reliability under Lawrence Berkeley National Laboratory Contract No. DE-AC02-05CH11231.

interpret results in the context of building physics to find important trends that explain their variability. We argue that control accuracy and baseline error drive inefficiency, and state-dependent dynamics create sensitivity to the DR event shape. We discuss the efficiency metrics used in the literature, the limitations of their application, and suggest an augmentation of an existing metric, which we call uninstructed energy consumption (UEC) to support evaluation of a wider, more realistic range of grid services. Finally, we explore these metrics with experimental data from Lawrence Berkeley Laboratory’s FLEXLAB. FLEXLAB is a unique facility with pairs of identical buildings that can be used for experimental applications, where each of the pair can serve as either a test or “control” building. Performing a new analysis on the FLEXLAB data from the frequency regulation experiments reported in [16], we show that the characteristics of the demand response trajectory have strong impact on the variability of efficiency metrics. We find that RTE can take on values comparable to the lowest reported in the literature if the trajectory is “asymmetric”, meaning that power requests on one side of the baseline exceed those on the other side of the baseline. We also illustrate that the UEC is much less sensitive to the characteristics of the demand response signal, providing more information about the building response’s inefficiency and inaccuracy by removing the energy content of the requested demand response service. We conclude that engaging buildings in demand response services need not result in large additional energy demands – as the RTE metric might suggest – but that control accuracy and the characteristics of the demand response signal are nonetheless important factors for a building operator to consider when evaluating the economics of providing demand response.

The remainder of the paper will be organized as follows: Section II describes the physical systems and mechanisms that we use to understand the behavior seen across the studies. Section III discusses each of the efficiency studies in detail and applies the physical mechanisms previously discussed to identify which may drive their results. Section IV discusses the use of the metrics of efficiency and proposes an additional metric, uninstructed energy consumption, to extend evaluation to asymmetrical DR services. Section V applies the metrics to data from previous experiments before final thoughts are presented in Concluding Remarks.

II. PRELIMINARIES

A. Variable Air Volume HVAC Systems

In the studies we discuss in this paper, demand response was provided by Variable Air Volume (VAV) HVAC systems that were cooling commercial office spaces. While all VAV HVAC systems designs are different, many share a common general architecture [24]. Commercial VAV systems typically have multiple zones controlled by independent thermostats throughout a building, and a single central plant that conditions and moves air. Figure 1 shows a diagram that depicts a representative HVAC system for these types of buildings. In these systems a chiller supplies chilled water to a cooling coil

that cools a mix of outside and return air from the conditioned space. The cooled air mixture is then distributed to all of the temperature controlled spaces via a variable speed fan. Since a single fan typically supplies air to a large number of zones in a building, each zone manages its own temperature by adjusting the position of a damper in its VAV box to manage conditioned air flow. Changes in zone-level damper position result in a change in the volume of conditioned air delivered to the zone. The central plant senses this increased demand and adjusts first its fan speed, and subsequently the thermal plant (typically just a chiller) output.




Fig. 1. General Variable Air Volume HVAC System Diagram.

The systems illustrated in Figure 1 typically have four feed-back control loops operating in tandem in order to maintain temperature in the spaces. The first feedback control loop regulates the temperature of water exiting the chiller; higher water temperatures entering the chiller will cause the chiller to work harder to return the water temperature to the regulated output temperature. The second feedback loop controls the valve position that regulates chilled water flow into the cooling coil in the air handling unit (AHU). This valve is adjusted in order maintain a constant temperature of the mixed outdoor and return air leaving the AHU. This flow control impacts the previous chiller loop by raising the chilled water return temperature, making the chiller work harder to cool it back to its setpoint. The third feedback control loop manages the speed of the air supply fan to maintain a constant duct pressure at the fan’s outlet. This third feedback loop will respond to pressure changes caused by changes in air flow through VAV boxes, and in turn will cause a change in the heat transfer across the cooling coil in the air handling unit, which will eventually impact the chiller consumption. The fourth feedback loop is governed by each zone thermostat; these regulate the damper on the associated VAV box to change the volume of cool air flow into the room to reach the desired local temperature setpoint. In the rare event that the cool air flow can not be reduced adequately to maintain the setpoint, a resistive reheat coil inside the VAV box may warm the air so that the space is not over-cooled. Of these control loops, the first and the third

loop directly impact electricity consumption (chiller power and fan speed controls), and the other feedback loops have largely an indirect effect.

B. Demand Response Applications

Demand response services that are centrally coordinated by a electrical utility or system operator typically fall into three categories. The first is demand curtailment during periods of high wholesale prices or peak load. The quantity curtailed is computed as the difference between an estimate of demand in the absence of a demand response action (a “baseline”), and measured power consumption. Building energy consumption is typically greater than the baseline following the curtailment event, but the total energy curtailed is often greater than the post-curtailment “rebound.” [25]

The second type of demand response is similar to the first, however the resource may be requested to increase demand as well as decrease. This type of flexible DR is of interest to areas with high renewable generation capacity that may wish to ensure wholesale electricity prices are positive, manage ramping conditions, minimize renewable generation curtailment, or manage local power quality conditions in the presence of high distributed generation. Grid services like these can also be measured from a baseline, and perturbations to the electricity consumer’s service, such as space heating, can result in changes to baseline electricity consumption after the DR performance period.

The last type of demand response services provide reserve-based ancillary services, such as synchronous reserve or frequency regulation. Frequency regulation, being the most lucrative of the reserve services in many markets [5], has received considerable attention from the demand response research community. A resource providing frequency regulation bids its capacity into a market and then follows an automatic generation control signal from the grid operator that is generally received every 4 seconds and includes a real power setpoint that the resource must achieve. In many markets, this is a symmetric service such that the reserve capacity is both an offer to increase or decrease the power of the resource in response to grid needs and is a roughly energy neutral signal.

C. Demand Response Impact on Energy Consumption

It is convenient for illustration to consider a DR signal that is “symmetric” in the sense that the requested power deviations above baseline are equal and opposite to the requested deviations below baseline. While most traditional DR and grid signals are asymmetric, this symmetric service ensures that the electrical energy delivered to the resource while providing the service is the same with an without the signal. In many cases, DR requires additional energy consumption relative to baseline in the period following this symmetric service provision to restore space temperature to the nominal value (we refer to this period as the “recovery” period). Citing the inefficiency of a battery as motivation, researchers have framed this additional building consumption as demand response inefficiency. To measure this inefficiency, Beil, Hiskens, and Backhaus [18]

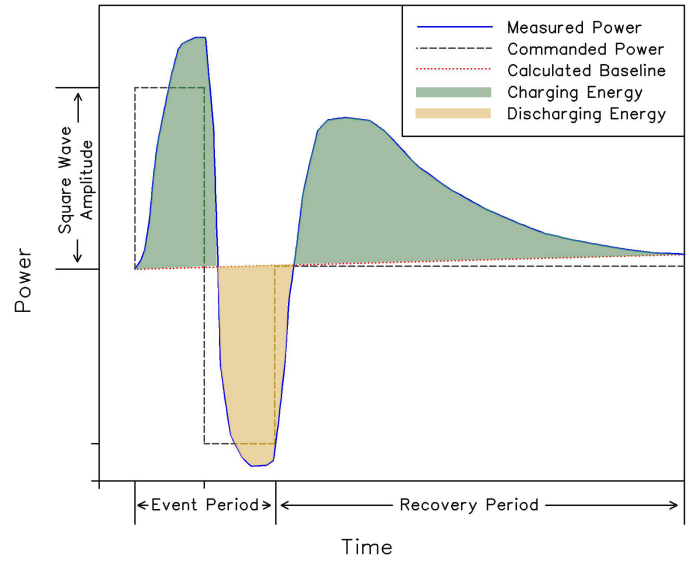


Fig. 2. An illustrative square wave demand response event. This image is based on Figure 4 from [18].

proposed illustrative “square wave” demand response events that mirror battery state of charge cycling. One of these events is a request for an increase in electrical demand followed by a symmetrical decrease in electrical demand, approximating a one cycle charge-discharge event of a battery, as shown in Fig. 2. Beil *et al* define any net increase or decrease in energy consumption relative to a baseline, during both the demand response and recovery periods, as the losses (or gains) to the system due to the square wave control.

Extending the battery analogy to HVAC-based demand response comes with an important caveat. In a battery, over a broad range of state of charge (SOC), the rate of change of SOC can be treated as independent of its current SOC state and linearly related to the input power. For HVAC, the best analogy for SOC is the temperature state of the conditioned space [21]. However, for HVAC-based demand response, the rate of change of temperature is strongly dependent on the temperature state. For example, if a building with a stable temperature in cooling mode undergoes a step increase in electrical power to the fan, the rate of change of the indoor air temperature will have a sudden initial increase, but slowly decline and eventually drop to zero as the space approaches a new stable equilibrium temperature. If the step change is held indefinitely, the temperature state will remain constant despite the increased electrical power input. In the battery context, electrical power input that has diminishing impacts on the state of charge would indicate a reduction of efficiency.

D. Efficiency Measures

The general question we address in this paper is: How much additional energy do buildings consume when they are used for demand response? There are two perspectives one can use to put this additional energy consumption into context. A “building centric” perspective is to measure a building’s

total energy consumption over a demand response period, including transient consumption after the requested service is provided, and compare this directly to an estimate of what the building would have consumed in the absence of the service (i.e., a baseline). The difference, defined as Additional Energy Consumption (AEC) in Keskar *et al* [20], will be positive and large for inefficient building impact. The ratio of AEC to baseline energy consumption provides a normalized efficiency penalty of the DR service. It is possible to have negative efficiency penalty, if the building consumes less energy during DR than in the baseline estimate.

Alternatively, a “service centric” measure of DR efficiency, originally proposed by Beil, Hiskens and Backhaus [18], measures energy curtailed or consumed for the DR service only. This involves partitioning building demand into periods when power is below baseline and periods above baseline. Invoking a battery analogy, we can regard the below-baseline energy as battery discharging, or E_{out} , and above-baseline energy as battery charging, or E_{in} . Continuing the battery analogy, the round-trip efficiency of the DR event is then simply the ratio of these, i.e. $RTE = E_{out}/E_{in}$.

E. Inefficiency Mechanisms

In the range of experiments we describe below, RTE is frequently below 1.0, often significantly so. This corresponds to large, positive AEC. What factors influence these differences in energy consumption? Here we consider a number of mechanisms that could lead to measuring RTE not equal to 1.0. The first four of these mechanisms are related to building physics and lead to real differences in energy consumption. The last two mechanisms, related to measurement and control, lead to differences in how we *quantify* efficiency.

State dependent losses. If the temperature of the building during a DR event is further from ambient conditions than during baseline, heat transfer will increase, and as a consequence more energy will be required to condition the space. Therefore if the building mass temperature is, on average, farther from ambient conditions during the demand response event than during a baseline, the building will consume more energy over the event than it would have in baseline, and the DR event will be perceived as inefficient. Note that the reverse is true as well – if average temperature over the DR event is closer to ambient, the DR service can be interpreted as having RTE greater than 1.0. The magnitude of this effect will be different for every building based on building construction parameters, such as insulation and thermal mass. Furthermore this mechanism does not derive from degraded conversion efficiency of the building HVAC system, but rather from the simple thermodynamic implications of operating the building at different temperatures.

A consequence of the state dependence of heat transfer is the bounded range for which the conditioned space’s temperature can take. Raman and Barooah [21] show that there exists a range that internal temperature may take that is bounded by the ambient temperature and the HVAC supply air temperature in most conditions. They use this fact to show that if the

square wave, or any symmetric DR signal, were repeated indefinitely, the efficiency RTE converges to one. This concept also has implications for the impacts of extending the duration of the symmetric DR event. As the length of the DR event increases, the closer the space gets to stable temperature states in each half of the charge-discharge cycle. When this happens, the VAV control response in the space during the recovery period becomes dependent solely on the stable equilibrium temperature in the second half of the DR event. The remainder of the losses will be accounted in changes to chiller energy consumption versus the baseline. Conversely, a very fast signal would be expected to have small impacts on the temperature of the space, and thus very little state dependent losses. Similar effects of service duration on energy consumption were seen in [23] in which faster services have little impact on energy consumption while maintaining temperature, but as duration increases so does energy consumption over baseline.

If state dependent losses dominate measured efficiency trends, we should expect to observe several phenomena. First, experiments and simulations with average temperatures closer to ambient should be more efficient than those with larger temperature differences. Second, holding all else equal, in cooling conditions increasing average zone temperatures should reduce energy consumption relative to the baseline, and vice versa. As we will discuss below, two major experimental studies do not meet these conditions, suggesting that state dependent losses are not the dominant explanation for measured DR inefficiency in those cases.

Nonlinear performance curves. The second factor that causes DR energy consumption to differ from baseline consumption is rooted in the relationship between HVAC equipment energy conversion efficiency and system loading. Very broadly speaking, HVAC chillers, pumps and fans are more efficient at part load than at full load. As a result, the relationship between system energy consumption and thermal energy loading – sometimes called the “performance curve” – is nonlinear. The characteristics of this efficiency-loading relationship mean that, holding all else equal, a system consumes more electricity if its output is variable than if it is constant, even if the net thermal transfer is the same in both cases.

If nonlinear performance curves drive measured DR efficiency trends, experiments with larger amplitude demand response trajectories should exhibit lower efficiency. Furthermore, experiments and simulations should always show efficiencies below 1.0. As we will discuss below, there is insufficient research to date that addresses the question of amplitude. However there are a number of papers – both experimental and simulation-based – that report efficiencies in excess of 1.0, suggesting nonlinear system performance is not the dominant factor.

Outside air mixing. VAV systems typically exhaust a fraction of return air and replace it with outside air to reject indoor air contaminants. The quantity of outside air introduced into supply air has an impact on chiller energy consumption as follows: if outside air temperature is greater than return air temperature, increasing outside air volume will increase

cooling energy required to maintain a particular supply air temperature. There are a number of ways to regulate outside air intake. At one extreme, the outside air *fraction* remains fixed across different supply air volumes, and at the other, the outside air *volume* remains fixed across different supply air volumes.

The impact of outside air intake on measured DR efficiency depends on how outside air fractions vary with total supply air volume. If the outside air *fraction* remains fixed, outside air would not cause the temperature of the air entering the cooling coils to vary with supply air volumes. Thus outside air mixing would have no impact on cooling energy required per unit of air delivered to the space, and there would be no impact on measured DR efficiency. If, on the other hand, outside air *volume* remains fixed, the temperature of the air entering the cooling coils would decrease with increasing supply air volume because proportionally less outside air would mix with the return air. In this case outside air mixing would tend to reduce cooling energy required per unit of air volume delivered as air volume increases. If none of the other mechanisms discussed here affect measured DR efficiency, constant outside air volumes would lead to a measured DR efficiency of greater than 1.0.

VAV reheat. In most VAV systems, ventilation air is provided by delivering conditioned air (which contains a fraction of outside air) to every zone. This leads to an undesirable side-effect: zones with low cooling loads will still receive cooled air in order to meet ventilation requirements, which could in turn over-cool those zones. To avoid over-cooling, VAV systems employ *reheat*, in which VAV boxes contain an auxiliary heat source to increase local supply air temperatures. Most reheat systems use hot water coils in the VAV box.¹ Should a significant fraction of zones be operating in reheat, the temperature response would derive from changes in hot water delivery rather than electrical power. Therefore in reheat conditions global temperature resets would result in smaller changes in electrical power than when zones are not in reheat. The resulting building performance is qualitatively the same as having a nonlinear performance curve: the mapping between electrical power and heat transfer is nonlinear, and more electrical power is required at high heat transfer rates. This effect should be more pronounced near minimum air flow at very low cooling loads. Therefore the qualitative observations we described above for nonlinear performance curves should also emerge if VAV reheat is driving inefficiency, particularly when building cooling loads are low.

Accuracy of control response. If a resource operates with an inaccurate and asymmetric, or biased, response to a symmetric signal, that bias will be recorded as inefficiency in the RTE metric and impact the resulting AEC. This control inaccuracy could be caused by poorly tuned parameters, excessive perturbations to the building, lack of feedback, or inaccurate control models. In any of these cases, the cause of

the poor efficiency metric is not an inefficient charge-discharge of the building, but rather is an artifact of the way in which efficiency is measured. For example, if the resource in Figure 2 were to only respond with half the power requested in the discharge period, the asymmetric and inaccurate response would halve the reported round trip efficiency despite the fact that the reduction in discharge energy was due to poor control and not due to conversion, thermal, or other energy losses. If this mechanism is driving reported inefficiencies in the literature, we would expect to see inaccurate and asymmetric control responses in experimental results, and results with greater observed inaccuracy and / or asymmetry would yield lower efficiencies.

Baseline calculation. Even small baseline errors can have an large impact on the measurement of RTE and AEC [19]. Baselines biased downward will increase measured energy for charging and decrease measured energy for discharging, which in turn would reduce the RTE. Similarly, upward-biased baselines would increase the RTE. One potential source of this bias is a difference between the *shape* of a building load profile and its baseline. For example, constructing a baseline by linearly interpolating between power measurements at the start and end of a demand response event will result in a baseline that is always less than a concave load profile, and always greater than a convex load profile. Load profiles for commercial buildings are generally concave down in shape during occupied hours [27].

III. STUDIES ON DEMAND RESPONSE EFFICIENCY

As we indicated in the introduction, there is a broad range of reported efficiency results for HVAC-based demand response. Round-trip energy efficiency varies across studies from below 40% up to around 130%, depending on the conditions of the test. In general, experimental results suggest much lower efficiencies than simulation studies. In this section we will detail the experiments and simulations presented in the literature [16], [18]–[21], discuss similarities across them and suggest the inefficiency mechanisms described in the previous section that may be present. While establishing which mechanism is dominant may not be feasible from the literature, we can rule out those mechanisms that are not consistent with the dominant patterns in their results.

A. Testing Conditions, Assumptions and Results

The present section describes critical features of the building environments, modeling assumptions, grid services provided, control methods, and metrics used in each study. It also relates these to the results reported. These conditions help identify potential causes for differing results among experiments and simulations. Table I describes differences among the simulations and experiments in the buildings and environments, modeled components and dynamics, characterization of the DR services provided, control approaches used, and even some variation in metrics used for determining results.

In Table I, the conditions can be related to loss mechanisms defined in Section II. The rows relating to building construc-

¹We will assume hot water is produced with a fuel other than electricity. Note also that some engineers recommend using electric resistance coils for reheat [26], but we will assume this is not the case in our analysis.

	Beil, Hiskens, and Backhaus [18]	Vrettos et al. [15], [16]	Lin et al. [19]	Raman and Barooah [21], [22]	Keskar et al. [20]
Type	Experiment	Experiment	Simulation	Simulation	Experiment
Building Size [m ²]	30,000	120	-	-	(3 Buildings) 9,100; 9,700; and 14,600
Building Loc.	Los Alamos, NM	Berkeley, CA	-	-	Ann Arbor, MI
# of VAV set-points	350	1	1	1	104, 193, and 109
DR Services Offered	Sq. Wave	Freq. Reg.	Sq. Wave and Freq. Reg.	Sq. Wave	Sq. Wave
Period of Square Wave [min]	30	-	30, 120	2-600	60
Relative Capacity Offer	5-10%	0-50%	10%, 20%, 40% [†]	20% [†]	NA [‡]
Fan Model	Linear [◇]	Cubic	Linear & Cubic	Quadratic	Linear [*]
Recovery Period [min]	90	-	< 300	Temp. Dependent	48,60
Baseline Method	Linear Approx.	Measured Control	Simulated	Simulated	Linear Approx.
Control Approach	Open-loop GTR	MPC & PI on Fan	PI on Fan	Ideal	Fixed GTR

[†] 20% of the baseline fan power, not the rated power of the installed fans.

[‡] No target power service offer, solely symmetric temperature adjustments.

[◇] Bounded, linear probabilistic relationship between VAV contribution to fan power and GTR, see [28].

^{*} Implied linear relationship between temperature setpoint and power.

TABLE I
TEST CONDITIONS FOR DR EFFICIENCY EVALUATIONS IN LITERATURE

tion and environment, as well as those describing the service being presented, all play a role in the state-dependent losses and effects. The rows defining the capacity offered, the type of service offered and the fan model employed can impact the degree to which non-linearity may effect system losses. Both the method for developing a baseline coupled with the recovery period used play a role in the error in measurement and RTE recorded that baselines can create. Finally, the control approach coupled with the number of setpoints the system controls plays a large role in the accuracy of response that may adversely impact the RTE results.

1) *Experiment: Beil, Hiskens and Backhaus:* Beil, Hiskens, and Backhaus ran experiments on a large commercial office building in New Mexico with hundreds of VAV boxes [18]. They use a square wave demand response profile to evaluate RTE. The square wave event was composed of an event period of 30 minutes, an amplitude of either 15 or 30 kW, and a recovery period of 90 minutes. The baseline is determined by linear interpolation between the measured fan power for a short duration immediately before a given DR event and the fan power immediately before the next DR event, at the end of the recovery period.

The authors employed an open-loop control approach that made global adjustments to all thermostat setpoints in the building together, called a global thermostat reset, to achieve model-based targeted power levels. This control approach was developed and used for this particular building in previous work to provide fast demand response services [28]. The authors use a system identification approach to approximate a bounded distribution of VAV contributions to fan electricity demand with a linear response to temperature setpoint changes. The open loop control approach was relatively inaccurate on the test building, with the response ranging from 70% to

120% of the commanded change in power for the best 60% of commands. It is possible that aggregate accuracy could improve if the control approach is deployed across a fleet of buildings [28].

Beil *et al* report an average round-trip energy efficiency of 46% on a total of 78 experiments. They ran 37 experiments in which the square wave performed a charging action then a discharging action, similar to the example in Figure 2. Though the paper does not report indoor air temperatures, we expect that these charge-discharge sequences would cause the average difference between indoor and outdoor air temperatures to grow over the duration of the experiment, relative to an uncontrolled baseline. These charge-discharge experiments resulted in average efficiencies of 61%. The paper also reports 41 discharge-charge experiments. In this case we expect the DR action to reduce the indoor-outdoor air temperature difference. These discharge-charge experiments resulted in an RTE of 0.34%.

Based on these results, we can reject the hypothesis that any one building physics effect is dominating Beil *et al*'s observed trends in measured DR efficiency. Specifically, as we explained in the prior section, if state dependent losses dominate results then an event beginning with a discharge will often result in efficiency *greater* than 1.0, events beginning with charging should have efficiencies less than 1.0 and less than discharge-first events. This qualitative behavior is not observed in Beil *et al*. Neither nonlinear performance curves nor VAV reheat are likely dominating the results, because no differences in RTE are reported between experiments with different amplitudes. Beil *et al*'s results also do not support a hypothesis that outside air mixing dominates the observed trends, since measured efficiencies far below 1.0 are observed in measured fan power alone and are not significantly changed

when the chiller consumption is included.

Beil *et al.*'s reported DR efficiency trends are, on the other hand, consistent with the possibility that measurement and control issues drive DR efficiency trends. Specifically, the results are consistent with the effects we would observe for a downward bias in the linear baseline due to the visible concave shape of the building load. However, it is not expected that this effect would play a significant role in wide differences in charge/discharge vs discharge/charge cycles. Beil *et al.*'s results are also consistent with the effects of inaccurate control. Figure 3, a reproduction of Figure 4 from [28] shows the model used for open loop control and the control error for a 2°F GTR. One can see a bias in the error – specifically, control response is on average higher than the model used for control – as well as larger error distributions for negative setpoint changes. The bias would tend to generate over response in charging energy and under response in discharging and thus lower reported RTE overall. The larger control error for positive changes in power (negative temperature setpoint changes) could also have bearing on the result that discharge-charge efficiencies are lower.

The proposition that control accuracy plays a large role in the unexpected efficiency outcomes is supported by a line in the caption for Figure 5 of Beil *et al.* In it they describe conditions of asymmetry in response arising from an initial request to discharge that does not occur if the initial request was to charge. If the system under performs during discharge, but charges at the full power request, this asymmetry results will read as an inefficiency when calculating RTE, though this may be mitigated to some degree by a change in response during the recovery period. This inaccurate control response applying most commonly to the discharge-charge cycle is consistent with the surprising difference in the reduced RTE of these types of events relative to those seen in other studies.

2) *Experiment: Vrettos et al.*: In our previous work, Vrettos *et al.* [15], [16], we demonstrated a hierarchical control approach for providing frequency regulation with commercial building HVAC, including evaluating the efficiency impacts of the control. The demonstration was performed at Lawrence Berkeley National Laboratory's Facility for Low Energy Experiments (FLEXLab) [29] in a pair of identical 60 m^2 test cells. Each cell is thermally isolated from its twin, has a dedicated air handling unit with its own variable speed fan, and manual dampers on the VAV Boxes. The thermal isolation and identical layout and construction of each pair enabled one cell to be treated as a measured baseline while the other cell performed the demand response service.

The hierarchical control used in [16] includes a day-ahead regulation capacity optimization, model predictive controller (MPC) to manage temperature via air flow, and a PI controller that tracks the regulation target in the fan's electricity consumption. This control replaced the traditional fan and VAV box control loops, but left the chiller and AHU control unchanged to manage supply air temperature.

The day-ahead optimization and MPC employ both a two-state resistive-capacitive (RC) model of the space and a cubic

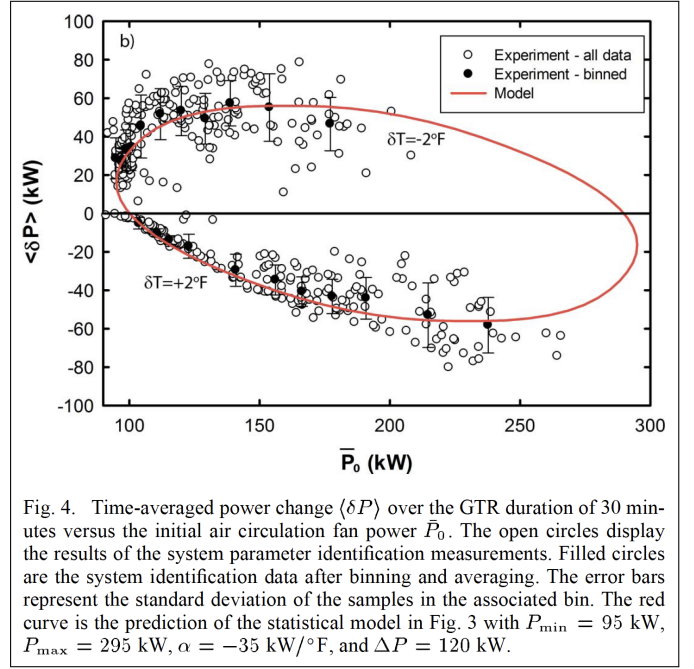


Fig. 4. Time-averaged power change $\langle \delta P \rangle$ over the GTR duration of 30 minutes versus the initial air circulation fan power \bar{P}_0 . The open circles display the results of the system parameter identification measurements. Filled circles are the system identification data after binning and averaging. The error bars represent the standard deviation of the samples in the associated bin. The red curve is the prediction of the statistical model in Fig. 3 with $P_{\min} = 95\text{ kW}$, $P_{\max} = 295\text{ kW}$, $\alpha = -35\text{ kW}/^\circ\text{F}$, and $\Delta P = 120\text{ kW}$.

Fig. 3. Experimental data and model for control approach used in [18], taken from Figure 4 of [28].

model of fan electricity consumption to manage the trade-off between the building's physical effects caused by state-dependent loss and non-linear performance and the economic value of offering the frequency regulation service. As the frequency of oscillations between charging and discharging is high when following the regulation service signal, the real-time tracking control would be expected to have very small impacts on the temperature state of the space and thus very little additional state-dependent losses. However, this highly variable signal could have an impact on non-linear performance losses due to the cubic fan response. The energy efficiency penalty from providing continuous service, calculated through both the energy delivered by the chiller to the spaces and the energy consumed by the fans, was reported as negligible, between 1-2% from the baseline cell, despite the variable signal. The low impact of inefficiency found for providing frequency regulation here is consistent with more recent experiments on heat pumps and rooftop units in highly controlled environments [30].

The other inefficiency mechanisms outlined above were not present: There was no VAV reheat, nor was outside air mixed into the supply air. Control accuracy was very high. Because of the test-control configuration, measurement error in baseline energy consumption was assumed to be low, and no linearization bias is introduced. This assumption of negligible baseline error was determined to be false during the analysis for this work and will be discussed in greater detail in Section V.

Insight into state dependent losses are possible due to the hierarchical control approach. The day ahead optimization gave the controller the ability to trade-off costs of state-dependent

losses and revenue from service provision. When comparing the system providing frequency reserves to a baseline in which energy minimization is the optimization objective, there was considerable additional energy consumption recorded: 68% more fan electricity and 11% more energy into the AHU. This is because the system must create headroom to provide symmetric frequency reserves in either direction, and so the system cannot operate at its temperature that minimizes state-dependent losses within a comfortable temperature range. This efficiency penalty has been discussed in other literature as an opportunity cost of providing reserves [17], but is not present in any of the other studies and so is not discussed as an inefficiency mechanism above.

3) *Simulation: Lin et al.*: Lin et al. [19] performed a simulation-based analysis of DR from VAV HVAC to improve the understanding of the drivers of inefficiency. They use a similar two-state RC model as [15] for the thermal dynamics of a space with a single aggregate building temperature zone. They ignore humidity and treat outdoor air temperature and internal heat gains as constant, resulting in a constant HVAC power demand when the system is at equilibrium. They assume a linear relationship between air flow and fan power for most scenarios tested on the grounds that VAV systems maintain constant pressure drop across a range of fan speeds. They do, however, run a scenario that includes a cubic fan model to evaluate its impacts. They also ignore reheat and outside air mixing. Their model uses PI control on the fan to manage room temperature setpoints in normal operation, and to manage fan electricity consumption during a DR event. Lin et al. perform a scenario-based simulation to evaluate impacts of adjusting multiple test conditions, including the magnitude of the DR signal, the duration of the DR event, reducing the controller gains, the capacitance of the space, and the ratio of outdoor air. The focus of their experiments are on the illustrative square wave DR service, but typically with a longer two hour event period and they record the system's post-event behavior until it returns to equilibrium.

Considering Lin et al.'s simulation setup, the only factors that could introduce measured inefficiency are control inaccuracy, state dependent losses, and baseline error. The results of their square wave experiments were highly consistent with a dominant mechanism of state-dependent losses. They reported RTE ranging from 0.81 to 0.94 for events that began with a charging event, and 1.05 to 1.19 for events that began with discharging. They also found that there were large impacts on RTE when changing event duration and by changing the building parameters to reduce available thermal storage. They also showed the impacts of control accuracy in tests in which the controller was tuned to be less aggressive, reducing overshoot and damped oscillation in the recovery period. The authors ran a test following a frequency regulation signal from PJM to compare to results found in [16], finding very little efficiency impact of providing the service which is consistent with the interpretation that state-based losses require enough event duration to see significant state change. Lastly, the authors explore the importance of baseline error, showing that

even a conservative error of 1% of the baseline load [31] can create an 8-10% change in RTE results.

4) *Simulation: Raman and Barooah*: In the studies by Raman and Barooah [21], [22], the authors propose formal definitions of a complete charge-discharge cycle, RTE, and the state of charge of an HVAC-based Virtual Energy Storage resource and explore RTE through theory and simulation. They define a complete charge-discharge cycle as one in which the state of charge at the end of the cycle is equivalent to its initial value, and that RTE is the ratio of the integral of Power out of the battery when discharging over power in during charging across the complete charge-discharge cycle. They also propose that state of charge for HVAC in buildings, given allowable temperature range, $[T_L, T_H]$ is the ratio of the difference between indoor temperature and the highest allowable indoor temperature, $T_H - T_r$, and the difference between the extremes of the allowable temperature range, $T_H - T_L$.

The authors go on to simulate square wave experiments using these definitions for RTE. The simulations use a single state RC thermal model, collapsing the room and the thermal mass into a single state. The model parameters were chosen to represent a large auditorium on University of Florida's campus. They use a quadratic model relating fan power and mass flow of the air for ease of analysis. Like [19], they assume constant outdoor air temperature and ignore humidity. The dynamic model is based on perturbing a system at equilibrium and allowing it to return. Control is idealized such that the square wave is perfectly tracked during the square wave DR event. To return the system to equilibrium control input is determined based on dynamic algebraic equations that relate the power and temperature dynamics to the equilibrium state before the DR event. The amplitude of the square wave is set to 20% of the equilibrium HVAC demand and the period of the wave is the parameter that is varied for analysis, between 2 minutes all the way up to 600 minutes.

Numerical results of the scenario-driven simulation for square waves that begin with charging show RTE ranging from approximate 93% for a very fast event down to roughly 72% for the longest period waves. In square wave simulations that begin with discharging, the opposite trend was observed: the very short event periods yielded the same roughly 93% efficiency but as the event periods increased the so did efficiency, up to around 128%.

The only modeled mechanisms that generate round-trip inefficiency are nonlinear performance curves and state dependent losses. The results indicate that the duration of the square wave dictates which of these loss mechanisms will dominate. For shorter event periods, non-linear losses in the fan dominate the RTE results, but as the event gets longer, the state-dependent losses become dominant as the average change in temperature in the space increases with the increasing duration. They show the existence of a critical duration in which the RTE of a discharge-charge DR event changes from being less than 1 to greater than one, and expresses it analytically based on parameters of the building thermal and fan model, the baseline operation conditions, and the amplitude of the square wave.

The authors also provide a mathematical proof that if the square wave is repeated n times, in the limit where $n \rightarrow \infty$ RT efficiency becomes one, which parallels the continuous service provided when offering frequency regulation from demand response. This result applies to any unconstrained repeating DR signal.

In [22], the authors extend their work looking at the continuous service provision case that results in 100% efficiency. They place a constraint on the temperature of the space such that the temperature over the course of the continually repeating DR event must be zero-mean. With this constraint, they show that there is a non-zero offset that must be applied to the square wave that results in an efficiency between 85% and 100%. The range of this result is driven by changing the building conditions and the duration of an event period. They conclude that non-linearity in both power and temperature dynamics drive the inefficiency under this temperature constraint.

5) *Experiment: Keskar et al.*: Keskar *et al* [20] ran more than 100 square wave DR tests across three buildings at the University of Michigan, Ann Arbor during the Summer months. These buildings ranged in size from approximately 9,100 to 14,600 m², with HVAC in each building controlling between 104 and 193 temperature setpoints. As with Beil *et al*, GTR was used, however in contrast to Beil *et al* there was no identification of a control law between a global temperature setpoint change and HVAC demand. Instead the control strategy implemented fixed, symmetrical temperature changes of either 2°F or 4°F. The paper assumed that a square wave *temperature profile* would result in a relatively square power profile. Each square wave experiment had a period of approximately one hour. The recovery period recorded for each experiment was either 48 minutes, due to a peculiarity of a single building's operation schedule, or 1 hour. Linear baselines were calculated by performing a least squares fit to the 5-minute period immediately before the event and the 5-minute period following the end of the recovery period.

The results for these experiments have a strong dependency on the size of the temperature setpoint change. The 4°F GTR experiments were consistent with Beil *et al*, with the charging-first experiments posting about 80% RTE and the discharging-first experiments around 48%. However, in the 2°F GTR experiments, the patterns were more consistent with Lin *et al*'s simulation results as 2 of the 3 buildings displayed significant *increases* in the RTE when the square wave was changed from charging first to discharging. These smaller GTR commands display behavior that is consistent with dominant state-based losses, whereas the larger GTR commands results are inconsistent with state dependent losses. The results due show sensitivity to amplitude, suggesting that non-linearity of performance and VAV reheat may play some role, but the drastic changes in the discharge-charge events between 2°F and 4°F suggest these are not the dominant mechanism. Control accuracy, on the other hand, cannot be ruled out as the dominant mechanism, especially for the large temperature setpoint changes. The study reports results for a total of 47 experiments, out of a total of 102 experiments

performed. The 55 “outliers” were filtered out due to poor control response (asymmetry in response exceeding 80% or no appreciable power response to control inputs), implying that the control approach used lacked accuracy and consistency. This inaccurate control coupled with results that are inconsistent with physics-based mechanisms of inefficiency suggest that inaccurate control may be the cause.

The AEC metric reported provides additional support for inaccurate control driving inefficiency. The experiment reports the AEC separately during the square wave event period and during recovery. If the GTR control was accurately providing a square wave response in power, the expectation would be that the AEC *during* the event period should be near zero. However the results indicate that roughly half of the overall AEC for each experiment was experienced during the event period, suggesting that the control response was highly asymmetric and inaccurate. This is corroborated in the graphical representations of events reported in the paper. Between the differences in the pattern of response to 4°F versus 2°F experiments, and the large amount of AEC present during the square wave event period, it is likely that the control inaccuracy is the dominant mechanism for the low RTE results for discharge-charge events.

B. Cross-study Comparison

The studies evaluating the efficiency of providing grid services with VAV-based HVAC Demand response provide a wide range of results. RTE as low as 34% and as high as 130% make overall evaluation of the efficiency of demand response challenging. However, the results presented also represents a large number of changing conditions, parameters and tests.

Table II presents the results from comparable conditions across the studies performing square wave experiments to examine the similarities and differences. The table shows the RTE results reported in each experiment for square wave with 30 minute event periods, the exception being Keskar *et al*. [20] which performed experiments for 60 minute event periods.

Table II shows considerable agreement across the square wave results that begin with charging. As all buildings and conditions for each experiment and simulation differ in construction and ambient conditions, we do not expect the results to be identical but that they are near one another. While both of the two simulation studies show losses driven by state-dependence with large differences caused by the order of the square wave and discharge-charge events with RTE greater than 1, comparing the two we see an expected reduction in efficiency in [21] as they modeled non-linearity in fan response. Both experimental studies record lower efficiencies than the simulations, which is likely due to significant bias due to linear baseline estimation of their concave down load profiles and possibly mediocre control performance during the event period.

There is a much larger disparity in results when evaluating the square wave experiments that begin with discharging. While simulation results are dominated by the state-dependence of thermal service delivery, resulting in RTE

Study	Round-trip Efficiency			
	Charge-Discharge		Discharge-Charge	
	η	n	η	n
Beil, Hiskens, and Backhaus [18]	0.61	37	0.34	46
Lin et al. [19]	0.94	-	1.06	-
Raman and Barooah [21]	0.88 [†]	-	1.01 [†]	-
Keskar et al. [20]	0.82 [‡]	24	0.76 [‡]	23
Keskar et al. [20] - 2°F Only	0.87 [‡]	11	1.08 [‡]	11
Keskar et al. [20] - 4°F Only	0.80 [‡]	13	0.48 [‡]	12

[†] Estimated from Fig. 3 in [21] with 30 minute square wave event period.

[‡] Average of results across all buildings from [20].

TABLE II
RTE RESULTS FOR SQUARE WAVE DR TESTS

greater than 100%, the experimental studies report results with considerably lower RTE than even the reversed square wave. The evidence strongly suggests that this disparity between simulation and experiment is driven by control inaccuracies that result in a lack of symmetry in the GTR control response. Evaluating the impact of control in Keskar *et al.* between the less aggressive 2 °F setpoint GTR and the 4 °F GTR, we see a marked difference. The 2 °F experiments result in average efficiencies close to the simulation results, however the 4 °F results mirror those in Beil, Hiskens, and Backhaus, in which discharge-first efficiencies are very low and less than charge-first. One mechanism that could lead to this effect might be running into minimum ventilation constraints. At equilibrium prior to the DR event, many of the VAV's in the system are too close to their minimum air flow position, and are unable to respond to the more aggressive temperature setpoint increase. However, in the second half of the experiment, there is ample capacity to respond to the increase. This would create an asymmetry in control response and a large bias towards inefficiency during the event period, well before the recovery period begins, which could explain these results. We expect that this or similar effects impacting control response explain the large difference in results between simulation and experiments, although a more detailed look into the responses and internal states of the building is necessary to confirm this hypothesis.

IV. EFFICIENCY METRICS AND APPLICATION TO DEMAND RESPONSE

The broad range of efficiency results could at least in part derive from RTE being an insufficient metric to capture the phenomena of interest. The present section discusses the battery analogy and state of charge in the context of efficiency. It goes on to discuss RTE and the metrics reliance on the symmetric square wave test. Finally, we discuss how AEC might be used to create other metrics that may have more meaningful and actionable information and apply to a larger variety of DR events.

A. The Battery Analogy, State of Charge, and Efficiency

As we discussed in Section II, the dynamics by which a battery changes its state of charge with an electrical power input are very different from that of an HVAC system. A

battery is generally approximated to have a linear dependence between its state of charge and power input. In the battery analogy, temperature is used as a proxy for the state of charge for thermal energy storage in an HVAC system. However, HVAC-based demand response has a strongly state-dependent relationship between the rate of change in temperature in the space and the heat removed from the space via cooling. In this section, we explore this dependency to highlight limitations of the battery analogy, namely that the SOC dynamics of HVAC are not proportional to power input like a battery and show that state dependence can cause the SOC to undergo greater change than a full cycle during a square wave event, which is recorded as a form of loss.

To clarify this state dependency, consider the simplified HVAC system represented by a single-state RC model used in Raman and Barooah [21]. For simplicity's sake, we will ignore internal heat gains so that the temperature dynamics described in [21] become:

$$C\dot{T}_r = \frac{1}{R}(T_{oa} - T_r) + \dot{m}_a C_{p,a} (T_s - T_r)$$

Where T_r is the room temperature. T_{oa} and T_s are the temperatures of the outside air, and the supply air, respectively. R and C are the resistance and capacitance parameters representing the buildings' insulation and thermal mass. $C_{p,a}$ is the specific heat capacity of air, and finally the \dot{m}_a is the mass flow of supply air. Qualitatively, the equation is a power balance that says the change in energy storage is equal to the the sum of the state dependent losses to the outside air and the heat removed from the space. The solution to the differential equation is:

$$T_r = \frac{aT_{oa} + \dot{m}_a b T_s}{a + \dot{m}_a b} + K_0 e^{-(a + \dot{m}_a b)t}$$

where K_0 is a constant determined by the initial temperature condition, $a = \frac{1}{RC}$, $b = \frac{C_{p,a}}{C}$, and we have assumed \dot{m}_a is independent of time.

Exploring the ODE and its solution gives three critical pieces of insight: First, the equilibrium temperature as $t \rightarrow \infty$ is a weighted average of the ambient outdoor air temperature and the supply temperature. Second, the exponential decay of temperature is dependent on the air mass flow rate. Third, temperature state dependence exists for both the losses to ambient and the thermal energy delivery from the HVAC, and their influence on the temperature profile is intertwined.

To further put this in perspective, consider an arbitrary square wave power command applied to both a battery and an HVAC unit. For a battery with a linear efficiency penalty placed on both charging and discharging, The symmetric square wave will result in a state of charge reduction proportional to the efficiency penalties and the magnitude of the square wave. The state of charge of an inefficient battery will always be less than its original state given a square wave input, regardless of charge/discharge order. Further, if charging is extended indefinitely, the battery must eventually reach 100% SOC and will stop charging.

As indicated by the solution to the ODE, HVAC temperature dynamics differ strongly from battery SOC when a square wave power command is imposed. The exponential decay term will cause the temperature to gradually reduce its rate of change when constant power is applied until the equilibrium temperature is reached. When a symmetric square wave in electrical fan power is applied to an HVAC system, the temperature will quickly change commensurate with the direction of the change in thermal power during the first half of the wave, but slowly decay as time continues. In the second half of the event, the change in thermal power delivery is larger than in the first, as the larger step change in supply air flow is compounded with temperature difference in the new state. This results in a more rapid change in the temperature than was observed in the first half of the wave. In fact, the temperature will arrive back at its original state before the period of the square wave is complete (assuming a near linear relationship between fan power and air flow) and the remainder of the time that the square wave is implemented will push the system past its original state of charge. Control will be necessary during recovery to make up for this excess charge, resulting in "losses" or "gains" that impact the RTE. An excellent example of this is visible in the default scenario simulation results in [19], reproduced here in Figure VI. Regardless of the order of the square wave, in the second half of the event the temperature crosses the baseline temperature, completing an SOC cycle, before the end of the event. After the event, PI control attempts to quickly restore the system to baseline by either increasing or decreasing fan power, causing the RTE reported to be above or below one depending on the order of the square wave pulses.




Fig. 4. Dynamics of HVAC System with a square wave fan input from Figure 4 in [19]

B. Service Definition and Round-Trip Efficiency

The present section explores the impact of service definition on RTE. We argue that RTE results for demand response from HVAC systems are heavily dependent on the service definition, which limits the utility of the metric.

In the experimental results summary of Section III we saw a strong dependence of the RTE results on the shape of the service requested. The order of the charge and discharge commands often results in a flip from efficiency gains to efficiency losses, due to state-dependent effects on heat transfer dynamics. This effect was magnified when increasing the duration.

When the concept of RTE was introduced, the service defined to generate results was a symmetric square wave. This illustrative service mimicked a simple charge-discharge cycle of a battery by asking for constant power in either direction with the same magnitude and duration. However, the standard definition for round-trip efficiency for a battery is the ratio of energy delivered in discharging to energy consumed in charging measured at the point of common coupling with the the grid [32], which is irrespective of the shape of the service. Dramatically different results can be possible if another service were defined.

Raman and Barooah [21] show that as a square wave tests are repeated N times and then allowed to recover after the repetitions, the RTE trends toward 1 as N grows. This is because perturbation to the temperature state due to any loss is bounded by the outdoor air temperature and the supply air temperature. Thus any control required to recover the system to its initial state during the recovery period is also bounded in terms of energy input, but the energy content of the square wave signal repeated N times is unbounded. So as $N \rightarrow \infty$, $E_{out}/E_{in} \rightarrow 1$.

Another example of the impact of service on RTE results is a demand response shed. Grid services provided by demand response are often shed events, in which electrical demand is reduced in response to a grid signal. The simplest shed event for VAV HVAC demand response would reduce fan power for a period of time, then return fan power to its initial value. In this type of event, assuming ambient conditions and occupancy are fixed, the temperature SOC would rise toward some new equilibrium temperature during the shed, and then would return to its initial value, given enough time, due to the state-based thermal dynamics of an HVAC system. From the perspective of the battery analogy, this return to the initial SOC would be accomplished without having to "charge" the system by increasing demand over baseline because the system was commanded back to its pre-event control state and given ample time for heat transfer to find its equilibrium, and thus would have an infinitely large RTE. Considering RTE for this simplified DR event provides no useful information and highlights the dependence of the metric on the service used to measure it.

Since RTE results exhibit this dependence on the shape of the service being provided, comparing resources based on

the metric is only valuable when they are compared for the same service. By extension, RTE results are useful when the service they are being calculated for is commonly used, which is not the case for the illustrated square wave service. The benefit of the RTE metric for batteries is that it allows the prospective resource owner to make financial decisions when deciding on which storage resource to purchase and how to operate it, as the RTE is relatively constant for most chemical storage use cases. It also provides some grid operators a parameter for scalable prediction of a resource's available capacity for market awards. However as discussed above, because RTE depends strongly on the shape of the grid service requested, it is less useful for these generic decisions for DR applications. RTE's strongest application for DR is to predict energy consumption, but because the metric varies based on environmental and service contexts, a more direct measure of energy consumption can provide the same quality of information with a simpler execution.

C. Alternate Metrics for Demand Response Efficiency

In this section, we discuss an alternate metric, AEC, proposed in the literature, describe its value over RTE, and propose a simple extension to compare asymmetric services.

Additional energy consumption can provide a useful metric for demand response resource owners or power systems operators where RTE fails. It represents the total quantity of energy consumption over the baseline across the DR event and recovery periods. Understanding the AEC caused by providing a demand response service can lead system operators to a better prediction of the net load and a more accurate dispatch. It can also allow the demand response resource owner to better understand their marginal cost of providing service, particularly if their energy costs. However, AEC alone is insufficient to inform future participation without some understanding of the context of services offered or provided.

Creating a unit-less quantity for comparison to other resources, such as efficiency, becomes less meaningful with the dependence of the response on both the service parameters and the states of the resource and its surroundings. AEC could be normalized by some derivative of energy content of the service, such as the γ metric described in Lin *et al.* [19] in which AEC is divided by the energy in one half of the square wave signal. With such a metric, the same resource can be expected to have a significant scatter of results based on the shape of that energy service and the conditions of the resource at the time of testing. In order to be most informative, a resource owner will need to perform simulations and/or experiments of their resource under varying conditions so that a probabilistic model of energy consumption, and thus cost, may be obtained and used to develop a market strategy.

Even with adequate data to support prediction using AEC, the metric was designed for symmetric services, such as the square wave, to highlight the asymmetry in response. However, grid services are rarely symmetric and when AEC is applied to an asymmetric service, the resulting value can be heavily

skewed by the power that was commanded by the system operator, obfuscating actual losses from providing the service.

We propose a simple adjustment to AEC to extend it to any asymmetrical service by subtracting the service's control signal from the resource's measured response. We call this metric the uninstructed energy consumption (UEC), and it is directly analogous to the metric by a similar name used in wholesale markets as the basis for uninstructed imbalance energy charges [33].

To formalize UEC, consider a time period, T , comprised of both the DR event and recovery period, such as the two hour experimental period used for each repetition of DR experiments in [18]. The AEC is the integral over T of the measured resource power, $P_m(t)$, less the baseline power, $P_b(t)$. The requested grid service signal, $P_{sig}(t)$, represents the sum of the baseline power $P_b(t)$ and a change from that power demand, $\Delta P_{sig}(t)$. We note that in some cases, system operators request $\Delta P_{sig}(t)$ from a baseline quantity, and in others they request $P_{sig}(t)$. The UEC is related to AEC by subtraction of what would be the change requested from baseline, ΔP_{sig} . Alternatively, this simplifies to the integral of the control error, irrespective of the baseline:

$$UEC = \int_T (P_m(t) - P_{sig}(t)) dt$$

If the DR service involves $\Delta P_{sig}(t)$ requests from the system operator, then knowledge of the building's baseline would still be required to construct this metric. On the other hand if the system operator requests $P_{sig}(t)$ directly, knowledge of the building's baseline power consumption is not needed. Baseline power consumption is always required for AEC.

Because UEC, like AEC and RTE, also does not consider the quantity of demand response capacity offered, or the shape of the system operator commands, it would be valuable to pair this data with data collected about the service provided to get a more holistic view of performance. UEC provides an ex-post understanding of the unintended costs of providing any grid service, and can be useful in conjunction with AEC to get a more complete picture of costs and opportunities for savings from improved control performance. Collecting historical results of UEC and coupling them with environmental and service conditions will allow the power system operator or DR resource owner to compose a probabilistic picture of the unintended energy cost of providing a service which can be useful in decision making. Inclusion of electricity and ancillary service price information would be necessary to make estimates of the cost of providing ancillary services.

V. METRICS AND EXPERIMENTAL RESULTS

In this section, we take a more granular look at data from the experiments documented in [15], [16] and perform a new analysis applying the three metrics described in this paper: RTE, AEC, and UEC.

Here we focus on frequency regulation experiments performed in LBNL's FLEXlab. We use two thermally identical test cells and designate one cell our "Experiment" cell in which

the HVAC system follows a frequency regulation control signal while the other is the “control” cell which generates our baseline. Figure 5 displays the frequency regulation offers, signal and response time series for the experiment in which symmetric frequency regulation was provided. The top plot shows the capacity offered and the incoming regulation signal around a zero baseline. The lower plot shows the baseline power, the regulation signal and fan response. The control methodology used for the control cell was that of a system that is “regulation-ready”, or is maintaining temperature and air flow states to hold regulation reserve capacity, but not actuating the fan to follow the independent system operator’s control signal. For three hours in the early morning, the day ahead optimization opted not to offer frequency regulation to maximize revenue in other hours, so the RTE, AEC, and UEC metrics for the fan are not reported for those hours.




Fig. 5. Symmetric regulation and reserve provision experiment.

Because the test cells at FLEXlab are built with identical layouts and components, we used a measured baseline electricity consumption for the fan in the control cell to determine the efficiency impacts in our original work [16]. The model predictive controller is used to establish the necessary air flow rate to maintain temperature and hold reserves for each 15 minute interval that defines the baseline electricity consumption. That air flow command is sent to both control and experiment cells, but is augmented with the frequency regulation signal response in the experiment cell. However, when evaluating data for the present analysis we found that fan consumption in the control cell was greater than consumption predicted by the fan model of the experiment cell learned from data. This difference in these baselines is illustrated in Figure 6, with the measured baseline in red and the calculated baseline used in this analysis in black.

Due to these observed differences in the electrical power response of the fans between the experiment test cell and the control test cell, a baseline derived from the measurement of the control cell’s fan power creates a bias in the results toward greater discharging energy. To avoid this bias, the

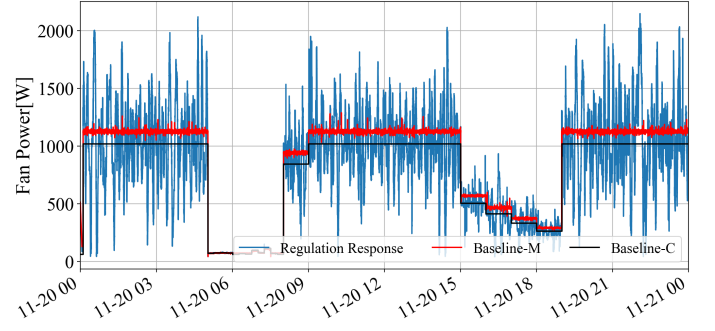


Fig. 6. Measured baseline challenges in symmetric regulation experiment.

baseline electricity consumption of the fan in this work is computed by running the air mass flow commands used to maintain the control room’s temperature (which are the output of the model predictive controller) through the cubic fan model obtained from data in the experimental test cell’s AHU. In this way, we preserve the temperature dynamics of the baseline conditions in the control test cell but ensure that the electricity consumption represented by the baseline is consistent with the machinery used to perform the regulation test.

Evaluation of the metrics of efficiency is performed on an hourly basis in this analysis to be consistent with the duration of a frequency regulation capacity offer in most US ancillary service markets. Table III provides statistics of RTE for the symmetric regulation experiment. The statistics it reports are the mean, standard deviation, median, min, max, 25th, and 75th percentile RTEs for the energy content of the frequency regulation signal, for the fan, an estimation of the chiller, and a combined result of the fan and chiller. As discussed in [15], the twin test cells share a single chiller, so it is necessary to disaggregate the impacts of each cells’ dynamics from total chiller power. In this case, we assume the same value for COP, 3.5, used in [19] and [21], and compute electrical power by dividing heat transferred across the coil by COP. Heat transferred across the coil is computed by the product of coil water mass flow, temperature change and the specific heat of water.

If the controls were operating with perfect accuracy, then the fan RTE would perfectly replicate the signal RTE. While some of the spread in the fan RTE can be attributed to the non-zero energy content of the regulation control signal, there is an average difference between the two in their hourly RTE score of 0.11, suggesting some inaccuracy in the control response. Once the chilled water loop was considered, the system’s RTE scores tended to increase with the majority of scores in the 0.9-1.1 range.

Table IV parallels the results shown in Table III, but for the day in which asymmetrical frequency regulation was allowed. In this relaxed condition, the system operates its baseline closer to the energy efficient optimum and offers considerably more regulation down (the ability to increase demand) than regulation up. The results quickly show the weakness of RTE to handle an asymmetric response. Both the fan and the signal

	Mean	Std. Dev.	Median	Min	Max	25th	75th
Signal	0.9874	0.1776	0.9674	0.7222	1.4238	0.8764	1.1461
Fan	0.8896	0.1519	0.8133	0.6603	1.1034	0.7723	1.0650
Chiller	1.1893	0.3920	1.0881	0.7373	2.2764	0.9383	1.2700
Combined	1.0472	0.2883	0.9801	0.6170	1.9848	0.8875	1.1425

TABLE III

STATISTICS OF HOURLY ROUND-TRIP EFFICIENCIES FOR SYMMETRIC FREQUENCY REGULATION EXPERIMENTS.

	Mean	Std. Dev.	Median	Min	Max	25th	75th
Signal	0.2524	0.1151	0.2223	0.1065	0.6242	0.1971	0.2733
Fan	0.2473	0.1135	0.2198	0.1039	0.6197	0.1898	0.2671
Chiller	0.9445	0.3768	0.8583	0.5525	2.0912	0.6982	1.068
Combined	0.7010	0.3540	0.6021	0.3568	1.7553	0.4838	0.7238

TABLE IV

STATISTICS OF HOURLY ROUND-TRIP EFFICIENCIES FOR ASYMMETRIC FREQUENCY REGULATION EXPERIMENTS.

results show very low RTEs, due to the greatly increased charging energy in the asymmetric signal over that of the discharging.

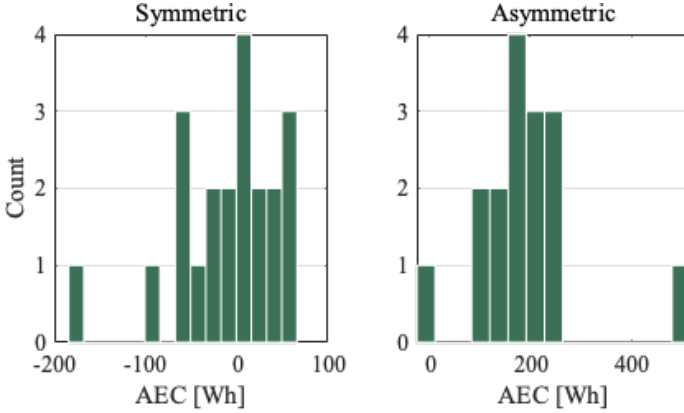


Fig. 7. Histogram of hourly additional energy consumption for both symmetric and asymmetric experiments.

Histograms of AEC for the two experiments are shown in the subplots of Figure 7. For the symmetric regulation experiments, AEC is averaging near 0 Wh and largely contained within -100 to 100 Wh. For the asymmetric case, almost all of the AEC recorded is greater than 100 Wh, getting as high as 600 Wh. This major difference in additional energy accounts for the energy inherent in the service that was offered, a cost or revenue that may already be expected, and doesn't represent any unanticipated losses or additional thermal energy stored in the space. Using AEC for asymmetric services offers provides limited understanding of any additional energy consumption that may be incurred by providing grid services and makes a performance comparison to days in which symmetric service was offered more challenging.

Figure 8 shows the uninstructed energy consumption that may result in a charge or credit that was not inherent to the regulation signal. The UEC results for both the asymmetric and symmetric cases are much more comparable than those of AEC: Both asymmetric and symmetric results in comparable magnitudes of uninstructed energy, though the asymmetric UECs are almost entirely positive (the capacity offered was

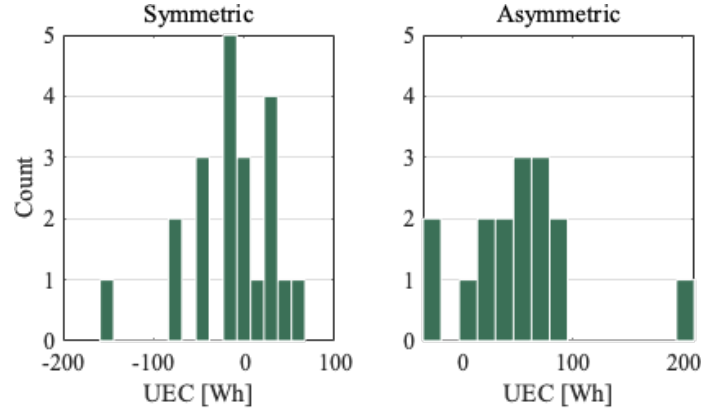


Fig. 8. Histogram of hourly uninstructed energy consumption for both symmetric and asymmetric experiments.

positive) while the symmetric case had both positive and negative UEC. This bias could provide insight into unexpected costs of providing asymmetric service or suggest inaccuracies present in the control that can be compensated for. In this case the largely positive skew of the UEC suggest that overshoot is dominating the control error. This insight is obfuscated in the AEC results by the presence of the energy in the asymmetric signal. While AEC still represents real costs and presents value to the resource owner, the advantages of using UEC allow for comparison of performance of a single resource across both symmetric and asymmetric services.

VI. CONCLUDING REMARKS

The literature on the efficiency of demand response presents a large – even confusing – array of results, including some round-trip efficiencies that are surprisingly low and suggest that large-scale reliance on demand response would result in significant increases in electricity consumption. When a square wave charge-discharge or discharge-charge cycle is commanded of VAV HVAC systems, simulation results suggest state-dependence of temperature dynamics will dictate whether the resulting RTE is less than or greater than 100%. However, experimental results display much lower RTE results than simulation, and appear to reverse the state-dependence effect. We lay out physics and engineering-based mechanisms that

could be leading to these results, then evaluate the available results in the literature to determine which mechanisms can be ruled out and which may be present. We propose that differences in RTE may be caused by bias introduced by linearization methods for determining the baselines, and the conflation of poor tracking response with efficiency loss and that actual efficiency losses may be much lower than the some of the experimental results indicate.

This paper also discusses the sensitivity of the RTE metric to the shape and definition of the grid service provided in the context of demand response. The logic behind the RTE metric stems from the battery analogy, but batteries and HVAC systems have very different responses in state of charge and temperature, respectively, to a constant electrical power control input. While batteries display a constant growth proportional to the input, HVAC systems have an exponentially decaying change in temperature that drives the system toward a new equilibrium. This vastly different dynamic response results in a strong relationship between the RTE results and the shape of the power commanded for grid services over time. This suggests that the RTE metric, as a means for comparison, is only useful to grid operators and resource owners if the grid service being evaluated is representative of a real service, which is not the case for the symmetric square wave.

While the additional energy consumption metric contains more actionable information for resource owners, it, like RTE, has limited applicability to asymmetric grid services, which are the most common class of services that DR provides. Both AEC and RTE have no means to distinguish between asymmetry in energy that was commanded versus asymmetry that stems from energy losses or inaccurate control response. The metric we propose – uninstructed energy consumption – uses the grid operators command signal as well as the resource’s control response to determine how much of the measured output was not requested by the system operator. This allows for the comparison of the impacts of providing varying degrees of asymmetric services with the same resource, which could provide the resource owner actionable information about control response in addition to evaluation of the unintended costs of providing grid services from HVAC-based demand response. In our experimental results, AEC values for symmetric services are 2-5 times less than those calculated for asymmetric services of comparable capacity, indicating a large bias in results caused by the asymmetry of the signal unrelated to inefficiency of service provision. In contrast, with the UEC metric, differences between symmetric and asymmetric performance are much smaller. Any remaining differences indicate true biases in unexpected energy in asymmetric service provision. These artifacts could provide useful information for a resource owner without obfuscation from the control signal.

However, these alternate measures of energy loss or gain still suffer from a dependence on the shape of the service, and would be most informative when a large dataset is gathered from simulation and/or experimentation with the intended HVAC system. Coupling this information with information

about the services provided and market prices could lead to valuable approaches for market participation.

The results of the present work suggest the need for the following research:

- Perform additional experiments and data analysis to isolate the dominant loss mechanisms and verify the impacts of baseline linearization and inaccurate control performance in discharge-charge experiments.
- Research to improve control methods, especially in large commercial buildings with many VAVs to control, is needed to help DR resource owners understand and contain their costs.
- Further development of strategies to use historical measures of AEC and UEC predict the energy cost of providing services – and to subsequently generate bid curves for market integration.
- A deeper understanding of the comfort implications of actuating demand response is needed, as these have the potential to be greater than the added electricity cost, particularly as the duration of grid service increases.

Research in these areas will support more informed market participation for future demand response resources.

REFERENCES

- [1] Y. Makarov, C. Loutan, Jian Ma, and P. de Mello, “Operational Impacts of Wind Generation on California Power Systems,” *IEEE Transactions on Power Systems*, vol. 24, no. 2, pp. 1039–1050, May 2009. [Online]. Available: <http://ieeexplore.ieee.org/document/4808228/>
- [2] J. Eyer and G. Corey, “Energy Storage for the Electricity Grid: Benefits and Market Potential Assessment,” Sandia National Laboratories, Tech. Rep. SAND2010-0815, Feb. 2010.
- [3] G. Strbac, M. Aunedi, I. Konstantelos, R. Moreira, F. Teng, R. Moreno, D. Pudjianto, A. Laguna, and P. Papadopoulos, “Opportunities for Energy Storage: Assessing Whole-System Economic Benefits of Energy Storage in Future Electricity Systems,” *IEEE Power and Energy Magazine*, vol. 15, no. 5, pp. 32–41, Sep. 2017. [Online]. Available: <http://ieeexplore.ieee.org/document/8011506/>
- [4] J. W. Zarnikau, “Demand participation in the restructured Electric Reliability Council of Texas market,” *Energy*, vol. 35, no. 4, pp. 1536–1543, 2010.
- [5] J. MacDonald, P. Cappers, D. Callaway, and S. Kiliccote, “Demand response providing ancillary services: A comparison of opportunities and challenges in the US wholesale markets,” Irving, TX, Dec. 2012, IBNL-5958E. [Online]. Available: <https://drcr.lbl.gov/publications/demand-response-providing-ancillary>
- [6] B. Foster, T. Bialecki, D. Burns, D. Kathan, M. P. Lee, and S. Peirovi, “2019 Assessment of Demand Response and Advanced Metering,” Federal Energy Regulatory Commission, Staff Report, Dec. 2019.
- [7] California Public Utilities Commission, “Order Instituting Rulemaking To Enhance the Role of Demand Response in Meeting the State’s Resource Planning Needs and Operational Requirements,” Sep. 2013. [Online]. Available: <http://docs.cpuc.ca.gov/PublishedDocs/Published/G000/M077/K151/77151993.PDF>
- [8] U. F. E. R. Commission, “Order 745 - Demand Response Compensation in the Organized Wholesale Power Markets,” Mar. 2011, docket No. RM10-17-000. [Online]. Available: <https://www.ferc.gov/EventCalendar/Files/20110315105757-RM10-17-000.pdf>
- [9] D. Callaway and I. Hiskens, “Achieving controllability of electric loads,” *Proceedings of the IEEE*, vol. 99, no. 1, pp. 184–199, 2011.
- [10] Y.-J. Kim, D. H. Blum, N. Xu, L. Su, and L. K. Norford, “Technologies and Magnitude of Ancillary Services Provided by Commercial Buildings,” *Proceedings of the IEEE*, vol. 104, no. 4, pp. 758–779, Apr. 2016. [Online]. Available: <http://ieeexplore.ieee.org/document/7416155/>

- [11] H. Hao, Y. Lin, A. S. Kowli, P. Barooah, and S. Meyn, "Ancillary Service to the Grid Through Control of Fans in Commercial Building HVAC Systems," *IEEE Transactions on Smart Grid*, vol. 5, no. 4, pp. 2066–2074, Jul. 2014. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6839081>
- [12] I. Beil, I. Hiskens, and S. Backhaus, "Frequency Regulation From Commercial Building HVAC Demand Response," *Proceedings of the IEEE*, vol. 104, no. 4, pp. 745–757, Apr. 2016. [Online]. Available: <http://ieeexplore.ieee.org/document/7407597/>
- [13] Y. Lin, P. Barooah, S. Meyn, and T. Middelkoop, "Demand side frequency regulation from commercial building HVAC systems: An experimental study," in *2015 American Control Conference (ACC)*. Chicago, IL, USA: IEEE, Jul. 2015, pp. 3019–3024. [Online]. Available: <http://ieeexplore.ieee.org/document/7171796/>
- [14] H. Yang, J. Zhang, J. Qiu, S. Zhang, M. Lai, and Z. Y. Dong, "A Practical Pricing Approach to Smart Grid Demand Response Based on Load Classification," *IEEE Transactions on Smart Grid*, vol. 9, no. 1, pp. 179–190, Jan. 2018. [Online]. Available: <http://ieeexplore.ieee.org/document/7442908/>
- [15] E. Vrettos, E. C. Kara, J. MacDonald, G. Andersson, and D. S. Callaway, "Experimental Demonstration of Frequency Regulation by Commercial Buildings – Part I: Modeling and Hierarchical Control Design," *IEEE Transactions on Smart Grid*, vol. 9, no. 4, pp. 3213 – 3223, Nov. 2016. [Online]. Available: <http://ieeexplore.ieee.org/document/7744683/>
- [16] —, "Experimental Demonstration of Frequency Regulation by Commercial Buildings – Part II: Results and Performance Evaluation," *IEEE Transactions on Smart Grid*, vol. 9, no. 4, pp. 3224–3234, Nov. 2016. [Online]. Available: <http://ieeexplore.ieee.org/document/7744657/>
- [17] D. H. Blum, T. Zakula, and L. K. Norford, "Opportunity Cost Quantification for Ancillary Services Provided by Heating, Ventilating, and Air-Conditioning Systems," *IEEE Transactions on Smart Grid*, vol. 8, no. 3, pp. 1264–1273, May 2017. [Online]. Available: <http://ieeexplore.ieee.org/document/7494593/>
- [18] I. Beil, I. Hiskens, and S. Backhaus, "Round-trip efficiency of fast demand response in a large commercial air conditioner," *Energy and Buildings*, vol. 97, pp. 47–55, Jun. 2015. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S0378778815002200>
- [19] Y. Lin, J. L. Mathieu, J. X. Johnson, I. A. Hiskens, and S. Backhaus, "Explaining inefficiencies in commercial buildings providing power system ancillary services," *Energy and Buildings*, Jul. 2017. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S0378778816320540>
- [20] A. Keskar, D. Anderson, J. X. Johnson, I. A. Hiskens, and J. L. Mathieu, "Do commercial buildings become less efficient when they provide grid ancillary services?" *Energy Efficiency*, Apr. 2019. [Online]. Available: <http://link.springer.com/10.1007/s12053-019-09787-x>
- [21] N. S. Raman and P. Barooah, "On the Round-Trip Efficiency of an HVAC-Based Virtual Battery," *IEEE Transactions on Smart Grid*, pp. 1–1, 2019. [Online]. Available: <https://ieeexplore.ieee.org/document/8738863/>
- [22] —, "Analysis of round-trip efficiency of an hvac-based virtual battery," in *5th International High Performance Buildings Conference at Purdue*, 2018.
- [23] Y. Lin, P. Barooah, and J. L. Mathieu, "Ancillary Services Through Demand Scheduling and Control of Commercial Buildings," *IEEE Transactions on Power Systems*, vol. 32, no. 1, pp. 186–197, Jan. 2017. [Online]. Available: <http://ieeexplore.ieee.org/document/7458167/>
- [24] J. D. Kelso, "Buildings Energy Data Book," U.S. Department of Energy, Buildings Technology Office, Tech. Rep., 2011. [Online]. Available: <https://openecr.org/doe-opendata/dataset/buildings-energy-data-book/resource/3edf59d2-32be-458b-bd4c-796b3e14bc65>
- [25] S. Kiliccote, M. A. Piette, J. L. Mathieu, and K. Parrish, "Findings from Seven Years of Field Performance Data for Automated Demand Response in Commercial Buildings," Pacific Grove, CA, Aug. 2010. [Online]. Available: <https://escholarship.org/uc/item/7m408394>
- [26] P. Raftery, A. Geronazzo, H. Cheng, and G. Paliaga, "Quantifying energy losses in hot water reheat systems," *Energy and Buildings*, vol. 179, pp. 183–199, Nov. 2018. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0378778818319467>
- [27] X. Luo, T. Hong, Y. Chen, and M. A. Piette, "Electric load shape benchmarking for small- and medium-sized commercial buildings," *Applied Energy*, vol. 204, pp. 715–725, Oct. 2017. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0306261917309819>
- [28] G. Goddard, J. Klose, and S. Backhaus, "Model Development and Identification for Fast Demand Response in Commercial HVAC Systems," *IEEE Transactions on Smart Grid*, vol. 5, no. 4, pp. 2084–2092, Jul. 2014. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6839121>
- [29] "Home | FLEXLAB," 2019. [Online]. Available: <https://flexlab.lbl.gov/>
- [30] J. Cai and J. E. Braun, "Laboratory-based assessment of HVAC equipment for power grid frequency regulation: Methods, regulation performance, economics, indoor comfort and energy efficiency," *Energy and Buildings*, vol. 185, pp. 148–161, Feb. 2019. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0378778818321571>
- [31] J. L. Mathieu, D. S. Callaway, and S. Kiliccote, "Variability in automated responses of commercial buildings and industrial facilities to dynamic electricity prices," *Energy and Buildings*, vol. 43, no. 12, pp. 3322–3330, Dec. 2011. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0378778811003707>
- [32] A. J. Wood, B. F. Wollenberg, and G. B. Sheblé, *Power generation, operation, and control*, third edition ed. Hoboken, New Jersey: Wiley-IEEE, 2013.
- [33] California Independent System Operator Corporation, "Fifth Replacement Tariff - Section 11: CAISO Settlements and Billing." Sep. 2019. [Online]. Available: <http://www.caiso.com/Documents/Section11-CAISOSettlements-Billing-asof-Sep28-2019.pdf>

AUTHOR BIOS



Jason MacDonald is a Principal Scientific Research Associate at Lawrence Berkeley National Laboratory, and a PhD Student in Energy and Resources at the University of California, Berkeley. He received an MS in Natural Resources and the Environment and an MSE in Mechanical Engineering from the University of Michigan, Ann Arbor in 2011. His research is centered on electricity market participation of distributed energy resources, with particular emphasis on the nexus of market interaction, optimization and controls. His work history includes engineering and research in plug-in electric vehicles, photovoltaic systems design, and demand response for ancillary services.



statistical learning and energy. His research focuses on grid integration of renewable electricity and models and control strategies for demand response, electric vehicles and electricity storage.

Evangelos Vrettos (S'09, M'16) is a Research and Digitalization Manager at Swissgrid AG. Prior to that, he was a Research Scientist at Lawrence Berkeley National Laboratory (Grid Integration Group) and then a Home Energy Systems Engineer at the Smart Grid Laboratory of the Swiss electric utility EKZ (Elektrizitätswerke des Kantons Zürich). He received the Dipl.-Ing degree in Electrical and Computer Engineering from the National Technical University of Athens (NTUA), Greece, in 2010. After his graduation, he worked as a research assistant at the Electric Machines and Power Electronics Laboratory of NTUA. In 2011, he joined the Power Systems Laboratory of ETH Zurich, where he obtained the Ph.D. degree in Electrical Engineering and Information Technology in 2016 and continued as a postdoctoral researcher before joining EKZ. His current research interests include demand-side management for renewable energy integration, optimal power flow applications for congestion management, and control of power grids with low rotational inertia.



Duncan Callaway is an Associate Professor of Energy and Resources at the University of California, Berkeley. He is also a faculty affiliate in Electrical Engineering and Computer Science, and a faculty scientist at Lawrence Berkeley Laboratory. He received his PhD from Cornell University. He has held engineering positions at Davis Energy Group and PowerLight Corporation, and academic positions at UC Davis, the University of Michigan and UC Berkeley. Duncan teaches courses on electric power systems and at the intersection of