INVITED PAPER

Frequency Regulation From Commercial Building HVAC Demand Response

The paper discusses the role of buildings for serving the grid by providing demand response (DR) and ancillary services. Commercial heating, ventilation, and air-conditioning (HVAC) loads are potential candidates for providing such DR services as they consume significant energy and because of the temporal flexibility offered by their inherent thermal inertia.

By Ian Beil, Ian Hiskens, *Fellow IEEE*, and Scott Backhaus

ABSTRACT | The expanding penetration of nondispatchable renewable resources within power system generation portfolios is motivating the development of demand-side strategies for balancing generation and load. Commercial heating, ventilation, and air conditioning (HVAC) loads are potential candidates for providing such demand-response (DR) services as they consume significant energy and because of the temporal flexibility offered by their inherent thermal inertia. Several ancillary services markets have recently opened up to participation by DR resources, provided they can satisfy certain performance metrics. We discuss different control strategies for providing frequency regulation DR from commercial HVAC systems and components, and compare performance results from experiments and simulation. We also present experimental results from a single ~30 000-m² office building and quantify the DR control performance using standardized performance criteria. Additionally, we evaluate the cost of delivering this service by comparing the energy consumed while providing DR against a counterfactual baseline.

KEYWORDS | Ancillary services; demand response (DR); frequency regulation; heating, ventilation, and air conditioning (HVAC)

Manuscript received January 5, 2015; revised November 12, 2015; accepted December 28, 2015. Date of publication February 16, 2016; date of current version March 17, 2016. The work at Los Alamos National Laboratory (LANL) was carried out under the auspices of the National Nuclear Security Administration of the U.S. Department of Energy under Contract DE-AC52-06NA25396. This work was supported by the Microgrid Program in the U.S. Department of Energy Office of Electricity. The work at the University of Michigan was supported in part by the National Science Foundation under Award CNS-1238962.

I. Beil was with the Electrical Engineering and Computer Science Department, University of Michigan, Ann Arbor, MI 48109 USA. He is now with Sargent & Lundy, LLC, Chicago, 60603 USA (e-mail: ianbeil@umich.edu).
I. Hiskens is with the Electrical Engineering and Computer Science Department, University of Michigan, Ann Arbor, MI 48109 USA.
S. Backhaus is with the Condensed Matter and Thermal Physics Group, Los Alamos National Laboratory, Los Alamos, NM 87545 USA.

Digital Object Identifier: 10.1109/JPROC.2016.2520640

I. INTRODUCTION

The proportion of wind and solar resources on the grid is rapidly increasing in response to energy policies that are encouraging less carbon-intensive generation portfolios. A heightened reliance on nondispatchable resources results in less available generation-side control, increasing the difficulty of operating power systems. To counter this trend, more emphasis is being placed on demand-side strategies to balance generation and consumption through adjustment of the load from its nominal value.

Demand response (DR) covers a broad class of demand-side control strategies that span wide time scales, magnitude and accuracy of response. Real-time pricing may be used to achieve DR [1]–[4], however, regulators have so far been reluctant to continuously expose retail customers to the volatility of wholesale electricity rates [5]. Instead, DR has typically been implemented using direct load control such as emergency load shedding [6], [7] or via long-term contracts that provide low energy prices for the right to curtail load [8]–[11]. Voluntary sustained DR is encouraged by fixed, time-of-use pricing or infrequent price increases during extreme system events, e.g., critical peak pricing [8].

Ancillary services markets are emerging as an alternative method to engage DR in power system control and as an additional source of compensation for customers capable of providing DR functionality. In particular, frequency regulation markets have opened up to nontraditional resources [12]–[15]. The rules governing these nascent markets are still in flux, but to qualify for participation, DR resources must pass performance benchmark tests and maintain acceptable performance scores over the long run [16].

Heating, ventilation, and air conditioning (HVAC) systems within commercial buildings are an attractive source of DR. As a significant portion of commercial building electricity consumption, they could create a considerable resource when aggregated across a sufficient number of buildings. Furthermore, a building's thermal mass buffers short-term fluctuations in net thermal load, enabling DR to perturb HVAC operation and its electrical consumption around nominal conditions with minimal impact on occupant comfort.

Subsequent discussions focus primarily on the application of HVAC for DR in large power systems. The ideas extend naturally to autonomous microgrids, where generation-demand balance is desired across isolated, relatively small connections of distributed generators and loads [17]. Campus-scale microgrids provide an example where building HVAC forms a large part of the total load [18]. In such a setting, generation-demand balance can be achieved through coordinated control of the HVAC loads of numerous buildings, other forms of DR, and generation resources.

This paper discusses different control strategies for providing frequency regulation DR from commercial HVAC systems and components, and compares performance results from our experiments and from experiments and simulations of other researchers. In particular, we consider the physical variables used to provide this control—direct control of fan speed or indirect control through manipulation of air mass flow rates, supply air pressure, or thermostat setpoints—and the tradeoffs between open-loop and closed-loop DR control.[1] To illustrate these ideas, we present experimental data from a $\sim$30 000-m^2 commercial office building where the HVAC electrical demand is modulated to track PJM qualification and historical frequency regulation signals. Performance of the HVAC system DR control is compared against standardized metrics [19]. We use these same experiments to quantify the excess energy required to provide frequency regulation service from this building.

The remainder of the paper is organized as follows. Section II provides an introduction to market-based frequency regulation from DR resources. Section III reviews typical commercial HVAC system architecture and the properties of several control strategies for providing frequency regulation from such HVAC systems. An overview of performance metrics is provided in Section IV. Section V reviews and compares experimental and simulation results from other researchers and from experiments performed by the authors. Section VI presents experimental results on the energy costs incurred while providing frequency regulation from our experimental testbed. Finally, Section VII offers a conclusion and suggestions for future work.

II. AN EMERGING DEMAND RESPONSE MARKET

In the United States, the Federal Energy Regulatory Commission (FERC) mandates independent system operators (ISOs) and regional transmission organizations (RTOs) to provide several ancillary services to ensure power system operability, including maintaining the system frequency at its nominal value. For significant frequency deviations, generating units respond locally through individual governor action [20], i.e., primary frequency regulation. Governor droop characteristics cannot restore nominal frequency, and centrally-controlled secondary frequency regulation [20], [21] is used for this purpose. Participating resources adjust their active-power setpoint according to a system-wide signal and are compensated through market mechanisms.

Secondary frequency regulation is typically provided by synchronous generators, but recent changes, most notably FERC Order 745[2] and FERC Order 755 [23], [24], have promoted demand-side resource participation. Together, these orders require that DR resources be fairly compensated in frequency regulation markets by utilizing performance-based metrics. Economic studies [25], [26] suggest that DR from the industrial and commercial sectors may soon be competitive with current ancillary service market clearing prices, although the situation varies significantly from market to market. A survey of the DR policy environment across various ISO/RTOs is provided by [13] and [14]. In one example, the Pennsylvania–New Jersey–Maryland (PJM) Interconnection has moved to accommodate increased DR into its frequency regulation market by providing a near-real-time market structure and allowing aggregation of a customer's demand assets [14].

Participation in the PJM demand-response frequency regulation market [16] (and similar markets) requires control actions on much shorter time scales and with higher tracking accuracy than previous DR applications. To enter and remain in these markets, participants must pass performance tests against standard benchmarks and

[1] Closed-loop control incorporates feedback that seeks to correct for discrepancies between actual and desired DR. In contrast, open-loop control has no mechanism for ensuring the actual DR tracks the desired response.

[2] As of writing, FERC Order 745 has been vacated by the U.S. Court of Appeals for the District of Columbia Circuit, although this decision has been stayed pending completion of a review by the U.S. Supreme Court, which presided over the case in October 2015 [22].

maintain acceptable metric-based operational performance. Section IV provides further details on these benchmarks and metrics, and examines the performance capabilities of a commercial HVAC system.

III. DEMAND-RESPONSE CONTROL IN COMMERCIAL HVAC APPLICATIONS

Buildings contain multiple electricity-consuming systems including lighting, computing loads, and security systems, as well as HVAC equipment. The last is uniquely qualified for DR participation because perturbations of an HVAC system (when well-designed) have a minimum impact on occupant comfort. Modulation of HVAC operation causes corresponding changes in the temperature of the occupied spaces, but the thermal mass of the building tends to buffer these effects, such that the temperature changes may be imperceptibly small (i.e., 1 °C–2 °C) to the building occupants. On the contrary, lighting systems provide a load that is adjustable at faster time scales and higher accuracies, but the flickering and outright darkness that would result make DR participation for lighting systems practically impossible.

Commercial and residential HVAC DR both have unique challenges. Residential HVAC units are controlled individually by simple hysteresis controllers that regulate building temperature within a specified deadband. DR requires coordination of large aggregations of these individually small thermostatically controlled loads (TCLs). Recent research has used simple first-order thermal models to describe this aggregate behavior and develop model-based control techniques [27]–[31].

In contrast, individual commercial HVAC systems (the focus of this paper) can provide a much larger controllable load than individual residential units. However, commercial HVAC systems have complex integrated subsystems, often involving diverse electricity-consuming equipment and interdependent control loops that create self-correcting behavior. Consequently, perturbing the electrical load of a commercial HVAC system to track a frequency regulation reference signal is challenging.

Previous experimental DR work using commercial HVAC systems focused on peak shaving applications and typically consisted of infrequent reductions in load sustained over multihour time scales; e.g., Piette *et al.* [32] demonstrated hour-long load shedding across geographically disparate commercial buildings under dynamic electrical pricing and Motegi *et al.* [33] utilized critical peak electricity pricing. This relatively infrequent but time-extended DR (load reductions only) can be accomplished by shutting down HVAC equipment or otherwise curtailing HVAC operations. However, frequency regulation service operates continuously at a subhourly time scale with both increases and decreases in load, and market performance metrics require a more accurate response than load curtailment. A better understanding of HVAC system architecture is needed to design a frequency regulation control system for commercial HVAC DR.

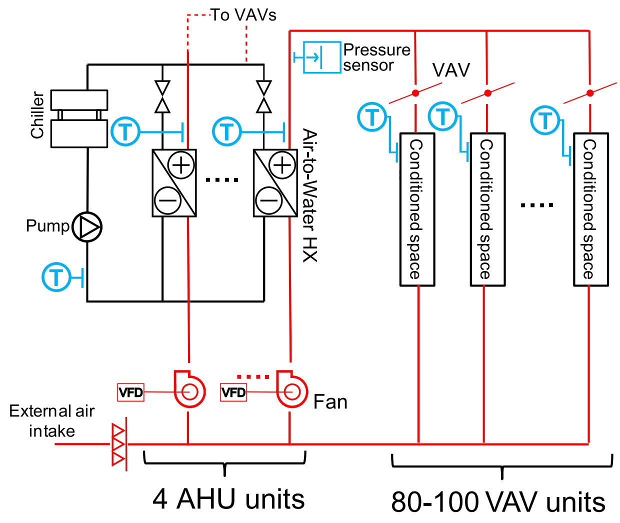

Fig. 1. ***A generalized HVAC system layout representative of the testbed used in our experiments [34]. It is composed of two main loops. The first is a water loop that circulates water between the chiller plant and air-to-water heat exchangers, where the water loop interfaces with the air loop. In the air loop, fans in the AHUs force warm return air from the conditioned spaces through the air-to-water heat exchangers where it is cooled and then delivered as supply air to the inlets of VAV units. Thermostats (T) in the conditioned spaces regulate VAV damper positions to control the amount of cooled air entering a conditioned space. Another thermostat on the chilled water loop regulates the chilled water flow to control the temperature of the cooled air supplied to the VAVs. A pressure sensor at the outlet of the AHUs regulates the fan speed to maintain a constant pressure for the supply air.***

A. Typical Commercial HVAC Architecture

Commercial building HVAC systems come in a variety of architectures with larger systems often being custom designed. However, their physical characteristics are typically similar to the architecture of the system used in our experiments [34] (see Fig. 1), and the experiments [35], [36] and simulations [35], [37] of other researchers. The HVAC system of interest consists of a central chiller plant that distributes chilled water to heat exchangers in several independent air handling units (AHUs). Each AHU contains a fan that circulates warm return air through the heat exchangers to supply cold air to the conditioned spaces. The flow of cold air into each space is regulated by a damper valve in a variable air volume (VAV) unit. Physically collocated VAVs are grouped together and connected via ductwork to a common AHU supply point (see Fig. 1). Our building testbed contains four AHUs, each serving ~100 VAVs. The supply air fan and associated variable frequency drive (VFD) in each AHU is controlled to generate a constant air pressure within the supply air duct.

Unlike the discrete, hysteretic control in residential HVAC units, the error signal from the conditioned-space thermostats is the input to a local proportional–integral–derivative (PID) controller that continuously varies the VAV damper valve opening between 100% (fully open) to about 20%–30% open [38]. The lower limit ensures the conditioned space always receives the required minimum level of ventilation. The local PID controller and the mechanical response time of the damper valve determine the rate at which the air flow responds to changes in the thermostat error signal. Typical response times are ~1 min in addition to any communication latency.

The heat removed from the warm return air is absorbed by the chilled water which is circulated back to the chiller plant. The chillers remove this heat and reject it to the ambient environment (see Fig. 1). There are two controls on the chilled water loop. First, water valves regulate the supply of chilled water to each air-to-water heat exchanger to control the AHU supply air temperature. A second controller regulates the chilled water outlet temperature through a two-level hierarchical structure. At the higher level, the controller adjusts the number of engaged chiller compressor units to achieve a coarse match to the necessary cooling load, while finer adjustments are made possible by managing the power of each compressor stage. Both of these controls operate on a time scale of ~10–15 min making them relatively slow compared to the VAV and supply air pressure controls.

The building automation system (BAS) provides communication and supervisory control for the entire HVAC system. Based on data gathered from a building's sensors and actuators, the BAS optimizes the operating setpoints for several key system parameters, including supply air pressure, supply air temperature, and chilled water supply temperature. The BAS updates these setpoints on a time scale of ~15–30 min, i.e., much slower than the VAV and AHU fan response times.

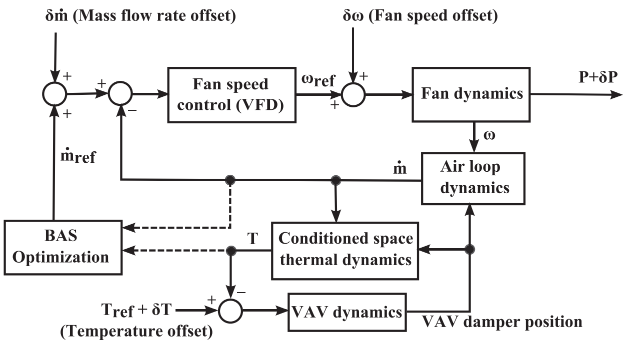

Fig. 2. *A simplified HVAC control diagram and the various control inputs that can be used to influence fan power consumption. The dashed lines indicate measurements that are telemetered to the BAS. Adapted from Fig. 4 [35].*

B. Commercial HVAC DR Control Methods

Modulating HVAC electrical load to track a DR reference signal over multiple time scales while maintaining occupant comfort is challenging due to the interdependent control loops within an HVAC architecture [39]. However, the focus on faster DR applications like frequency regulation narrows the potential control options. For example, infrequent chiller control for peak shaving has been thoroughly examined [40]–[42], but their slow response and potentially high on/off cycling make them inappropriate for frequency regulation. The remainder of this section focuses on leveraging the faster responding AHU fans for frequency regulation.

Fig. 2 schematically shows the multiple control loops that affect fan power. These loops can be modified in several ways to influence fan electrical load. Here, we focus on three methods:

- fan speed offset—directly adjusting fan speed through control of the VFD, e.g., by adding a fan speed offset signal $\delta\omega$;
- supply pressure/mass flow setpoint offset—adjusting supply mass flow (or supply pressure) setpoint by adding an offset $\delta\dot{m}$ (or δp) which then modifies fan speed via a local control loop that regulates supply pressure or mass flow;
- thermostat setpoint offset—adjusting the thermostat setpoints to modify the VAV opening via the local PID controller with subsequent impact on air flow and fan speed via cascading effects through the supply pressure/mass flow control loop.

Each of these control strategies may be implemented in an open-loop or closed-loop configuration. They are concisely summarized in Table I, with Section V-C providing a comparison of their technical performance. Here, we give a brief description of the control method and comment on some qualitative features. We refer the reader to the original references for detailed descriptions.

1) Fan Speed Offset: The fan speed offset method is the most direct way to influence fan power consumption and has been implemented in open-loop [36] and closed-loop [35] configurations. Both implementations require some level of system identification, either an experimentally determined transfer function [35] or a trained predictive model for power changes [36]. In either case, an offset is added to the motor speed signal to modulate the fan motor's power consumption. The number of supply fans in a commercial HVAC system is typically small creating some advantages. First, it limits communications and associated latency. Second, the VAVs are not directly involved in the control which avoids their mechanical latency. Finally, the control input is as close to the power consuming load as possible reducing the

Table 1 Summary of HVAC DR Control Methods.

Control method	Offset signal	Response rate	Tracking accuracy	Disruption to HVAC operation	Controllable duration	Approx. control bandwidth (cycles/min)
Fan speed offset	$\delta\omega$	Fastest	Highest	Moderate to high	Sub-hourly	0.1–1.0
Supply pressure/mass flow offset	δp or $\delta\dot{m}$	Fast	High	Moderate to high	Sub-hourly	0.1–1.0
Thermostat set-point offset	δT	Slow	Moderate	Minimal	Long term	0.02–0.2

complexity and uncertainties, and likely improving DR reference signal tracking. In fact, this form of control has been used to mitigate photovoltaic generation variations [43] which are faster than the frequency regulation signals considered here.

There are some drawbacks to supply fan speed offsets though. It may require retrofits to legacy VFD hardware adding cost and complexity. Also, downstream control loops regulating supply pressure or mass flow, or VAVs controlling conditioned space temperature, will compensate for changes to supply fan speed, limiting the ability to track reference signals with time scales longer than approximately one minute. While this effect can be mitigated in the short term (30–40-min intervals) by filtering out the low frequencies of the DR input signal [35], longer term (> 1 h) experiments using fan speed control have not yet been demonstrated. Motor speed ramp-rate limiters in the VFD likely limit the highest response frequency. When controlling fans that supply many VAVs, e.g., $\sim$100 as in [34], this method lacks the ability to customize DR control for specific occupants or conditioned spaces that may be particularly sensitive to HVAC variability.

2) Supply Pressure or Mass Flow Setpoint Offset: In many respects, supply pressure or mass flow setpoint offsets are very similar to supply fan speed offsets. Whether supply pressure or mass flow is used depends more on the design of the HVAC control system and is not really a choice of the DR control designer. This method has been implemented in closed-loop form in [35], and an open-loop form has been studied in simulation [37]. In either case, the setpoint offset to the respective control loop forces that loop to modify its input to the motor VFD, which ultimately changes the fan motor speed and power consumption. The advantages of this method are similar to the supply fan speed offset method, i.e., low latency from a small number of endpoints and no purposeful involvement of VAVs. The complexity is still relatively low, but somewhat higher than supply fan speed offset because of the physics of the supply duct work and the response of the pressure or mass flow controller. However, the physics of the flow in the duct can be explicitly incorporated [37] enabling accurate DR reference tracking. Finally, only software changes are anticipated because the supply pressure or mass flow setpoints can generally be modified via the BAS.

Drawbacks to this method are also similar to supply fan speed offsets. To maintain conditioned space temperature, the downstream VAVs will compensate for the imposed setpoint changes limiting the ability to track DR reference signals over longer time scales [37]. The response of the local pressure or mass flow control loop will likely limit the highest response frequency. When controlling the supply pressure or mass flow in an AHU that supplies many VAVs, e.g., $\sim$100 as in [34], this method has the same limitations as mentioned for fan speed offset.

3) Thermostat Setpoint Offset: Thermostat setpoint offsets [also termed global thermostat reset (GTR)] are the least direct control method of the three. This method has been implemented in open-loop form in [34] and in the present work, and has been simulated in [37]. We do not know of an implementation or simulation study of a closed-loop form. If the thermostats in a group of VAVs are adjusted to a cooler setpoint, local PID controllers on the VAVs open their damper valves to increase air flow. The subsequent drop in upstream supply air pressure forces the supply pressure control loop (discussed above) to increase supply fan speed (via the VFD) resulting in higher electrical power consumption. Similarly, a decrease in electrical power occurs when thermostats are adjusted to a warmer setpoint.

Advantages and disadvantages of thermostat setpoint offsets are mostly reversed relative to the previous two methods. Communicating with every VAV in a large building can create significant latency ($\sim$1 min in our testbed [34]) and the direct involvement of the VAVs in the control adds their mechanical latency of $\sim$30 s [34], [37]. The control is now quite complex as it relies on the behavior of a large number of VAVs, conditioned spaces and occupants, all of which are subject to many disturbances and not easily or accurately modeled. Statistical models have been developed to predict this behavior, but their accuracy in experiments is limited [34]. Simplified thermohydraulic models of the building and HVAC system have been used in simulation[37], but these were not subjected to significant disturbances or changes in nominal conditions.

This control method also has some significant advantages. Only software modifications are needed because remote control of the thermostat setpoints is possible through many BASs. This also enables tailoring the participation to individual VAVs with the ability to exclude particularly sensitive conditioned spaces or occupants from DR control. We note that this method enables simple (and easily verified) guarantees on occupant comfort through limiting the range of thermostat offsets. Although the initial response is slow because of latency, thermostat setpoint offsets are least susceptible to HVAC control system self-correcting effects and can track DR reference signals with longer time scales. In [34], tracking of square-wave reference signals with 30-min steps was achieved with no significant rolloff of the response.

IV. PERFORMANCE METRICS

FERC Order 755 [24] stipulates that DR resources must be compensated for the quality of service they supply to the electrical grid, with the specifics of the quality assessment left to each ISO/RTO. To improve the assessment of quality, the PJM demand response frequency regulation market divides the frequency regulation burden into two components: traditional regulation (termed RegA) and dynamic regulation (RegD) [19]. RegA is a low-pass-filtered area control error (ACE) signal[3] designed for ramp-limited DR resources that cannot adjust their demand quickly. RegD is a high-pass-filtered ACE signal for fast ramping resources that are capacity limited, e.g., flywheels and batteries. A more thorough discussion of the PJM market, and RegA versus RegD, is given in [37]. Resources bidding into the PJM regulation market specify the total up and down regulation that they can provide, e.g., $\pm$120 kW. PJM calls upon these resources by broadcasting a signal in the range $[-1, 1]$which the resources locally scale to their cleared capacity.

Before a DR resource can participate in the PJM market, it must demonstrate adequate performance against a standard 40-min qualification reference signal that commands the DR resource to adjust its load over its full capacity range, as shown in Fig. 4. Performance scores for both RegA and RegD participants are generated based on three criteria listed in [16], which gives a detailed description of each component. The following is a brief summary of the component score calculations.

- Delay: The time delay (rounded to the nearest 10 s) that provides the maximum correlation between the reference signal and the DR response. The score linearly decreases with delay, where a delay of 10 s or less nets a perfect 1.0 and a delay of 5 min 10 s or more scores a 0.
- Correlation: The maximum correlation between the reference signal and the DR response with the measurements time shifted to remove the delay calculated above. Although correlation may range over $[-1, 1]$, the correlation score is restricted to the range $[0, 1]$. Hence a negatively correlated relationship between signal and (time-shifted) response receives a score of 0.
- Precision: The absolute value of the error between the reference signal and the DR response, normalized by the average reference signal, is calculated for each 10 s interval. The precision score is given by 1 minus the average of the normalized errors. Unlike the correlation score, this component does not compensate for latency in the response, and hence a time-delayed response will impact not only the delay score but also the precision score.

The aggregate score is a weighted average of these three components (currently PJM weights each component equally [16]). To qualify, DR resources must pass three qualifying tests with an aggregate score of 0.75 or better. They must maintain a 100-h rolling average aggregate score of 0.40 or better to remain in the market.

V. EXPERIMENTAL TESTS

Qualitative advantages and disadvantages of the different control methods were discussed in Section III-B. In this section, we present details on the experiments carried out on our testbed building using PJM RegA qualification test signals, and compare these results to experimental and simulation studies utilizing different control methods. We then examine the behavior of the testbed building when

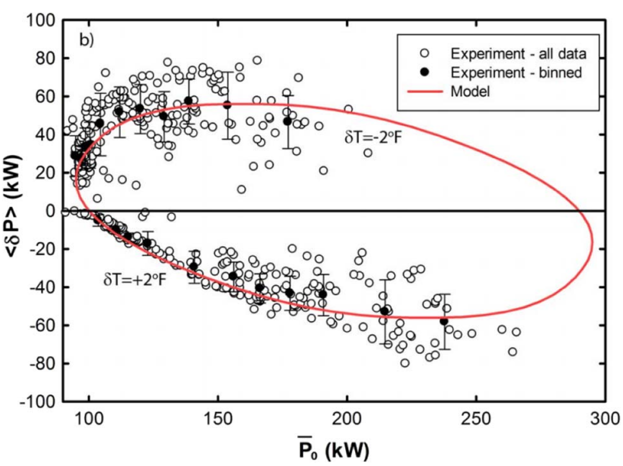

Fig. 3. ***A reprint of Figure 4 from [34], which displays a reduced-order model (the red ellipse) that relates initial fan load and the resulting change in fan power when both*** $+2\,^{\circ}$F ***and*** $-2\,^{\circ}$F ***thermostat offsets are applied. The same testbed building is used in both [34] and the studies reported in Section V. Reprinted, with permission, from [34].***

[3]The ACE signal for a balancing authority (control area) is a weighted sum of the mismatch between nominal and actual system frequency together with the mismatch between scheduled and actual power flows on tie lines to adjacent balancing authorities [20], [21].

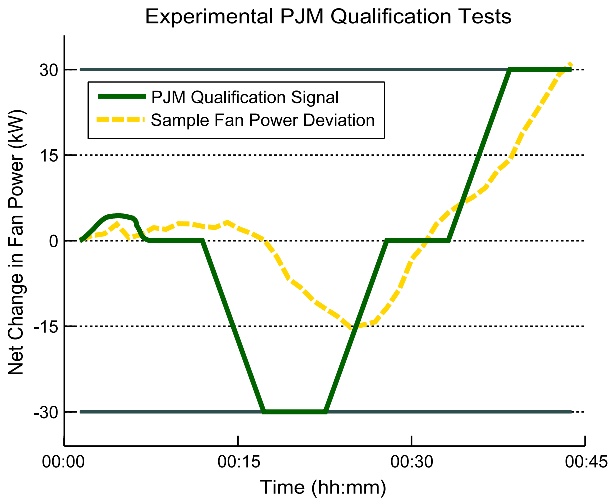

Fig. 4. ***Experimental results from our 30 000-m² test building. The 40-min PJM RegA qualification test (solid green line) is broadcast to the BAS. Using the control method developed in [34], the HVAC system uses the thermostat setpoint offset method to adjust fan power to track the reference signal. An example of the systems ability to track the regulation signal is illustrated by the yellow dashed line showing fan power deviation from nominal.***

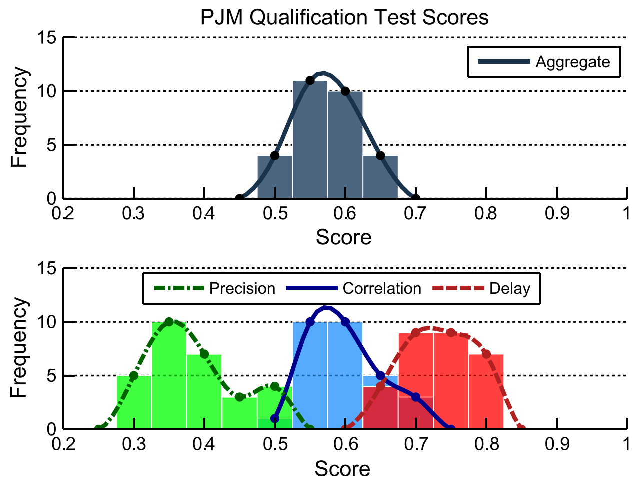

Fig. 5. ***Histograms of the PJM performance scores for the tests using the 40-min PJM RegA qualification test signal. Upper plot: Aggregate scores. Lower plot: Delay (red), Correlation (blue), and Precision (green) scores.***

introduced to historical RegA market signals, which are less aggressive than the RegA qualification test.

A. Control Scheme

Th DR control in our testbed building uses an open-loop thermostat setpoint offset method. This control law is based on the relationship, shown in Fig. 3, between changes in thermostat setpoint and the corresponding changes in fan power. A complete discussion of these experimental results is provided in [34]. The reduced order model shown as the red ellipse in Fig. 3 provides an (approximate) mapping from a desired change in fan power to the corresponding thermostat temperature offset. Tracking a DR reference signal involves mapping the required changes in the fan power to the corresponding thermostat setpoint offset which is then telemetered to the VAVs.

B. RegA Qualification Test

The qualification test experiments were conducted over several weeks in summer and early fall 2014, when HVAC load is at its peak. Six 2-h tests ran from 07:00:00 to 19:00:00 each day consisting of a 40-min period in which thermostat setpoints were adjusted to track the PJM RegA qualification test reference signal, followed by an 80-min period in which control was released so that the system could return to nominal operation. The regulation signal was scaled to ± 30 kW, which is not large enough to meet the 100-kW minimum for participation in the PJM market,[4] but is sufficient to induce observable nonlinear behavior in the HVAC system. An example of the experimental fan response is shown in Fig. 4, and a histogram of all the qualification test scores is presented in Fig. 5.

The observed aggregate scores for this building and control system are in the 0.5–0.65 range and fall short of the 0.75 score required for qualification. Each component score can be analyzed in light of limitations of the control methodology, communications and the particular HVAC equipment involved. The average delay score of $\sim$0.75 implies a time delay of about 1.5 min which is consistent with the typical combined communications latency ($\sim$1 min) and VAV response time ($\sim$30 s) for this testbed. The delay score could be improved in several ways. Other BASs use communications protocols with lower latency than the serial protocol used in our testbed and so should perform better. However, modifying existing BAS communications just to improve DR control for frequency regulation is not likely to be cost effective. A BAS with slow communications might be better suited for direct fan control or supply pressure/mass flow setpoint offset control schemes. More sophisticated control may also improve the delay score. For instance, model predictive control and state estimation are used in [44] to improve DR tracking performance of residential HVAC systems connected by a constrained communication network.

[4]PJM requires the total capacity of a resource to be $\geq$100 kW. The testbed building has a maximum nameplate fan load of 295 kW, although it typically operates closer to the minimum fan power of 120 kW (determined by ventilation requirements). At this lower power rating, a $\pm$30-kW change already represents a $\pm$25% change in fan power, and hence a 100-kW adjustment is infeasible for the individual building. Aggregating several building loads would be one way to achieve the necessary DR capacity.

Table 2 Summary of PJM Performance Scores

Control Method	Test	Agg	Del	Corr	Pre	Source
Fan speed	Sim	0.89	-	-	-	[35][1]
	Exp	0.77	1.00	0.83	0.49	[35][2]
	Exp	-	-	$\sim$0.90	-	[36]
Supply pressure/ Mass flow rate	Sim	0.77	1.00	0.87	0.44	[37][3]
	Sim	0.82	-	-	-	[35][4]
	Exp	0.81	0.95	0.98	0.50	[35][5]
T-stat set-point	Sim	0.82	0.99	0.92	0.54	[37][6]
	Exp	0.57	0.73	0.60	0.38	LANL[7]

[1] Table I, ASHFS (True output), Closed-loop, Reference=Fast ACE
[2] Table IV, ASHFS, Closed-loop, Reference=Fast ACE
[3] Figure 5, Open-loop, Reference=RegD
[4] Table I, ASLAF (True output), Closed-loop, Reference=Slow ACE
[5] Table IV, ASLAF, Closed-loop, Reference=Slow ACE
[6] Figure 8, Open-loop, Reference=RegD
[7] Current work (Figure 4), Open-loop Reference=RegA

The correlation scores of 0.55–0.60 are likely limited by the complexity and open-loop nature of the thermostat setpoint offset control methodology. Although this control methodology provides for easy customization to individual building occupants (e.g., we excluded $\sim$10% of the VAVs based on perceived sensitivity) and simple power baseline tracking, it is probably the least accurate in tracking a frequency regulation reference signal. It cannot correct for uncertainties and exogenous disturbances to the HVAC system caused by ambient environment changes, occupant level variability or other disturbances. However, we note that online identification techniques [45] could be used to regularly update the control law and potentially improve the correlation score.

The low precision scores (0.30–0.50) are caused by both time delays and inaccuracies in the open-loop scheme. The relatively fast ramping of the RegA qualification test signal compared to the communication latency creates control errors during ramping that are somewhat compensated by the time shift in the correlation metric, but create significant penalties in the precision score.

C. Comparison of Control Methods

The experimental results from Section V-B are listed in Table II (the final entry in the table, labeled "LANL"), along with available experimental and simulation studies,[5] to compare the DR control methods based on their expected technical performance. We restrict the range of studies to those focused on providing frequency regulation and that have used reasonably similar test procedures, i.e., the use of PJM metrics [16].

[5]The entries in the footnotes of Table II refer to the original references and specify the source of the results, the name or acronym for the control method, whether the control is open or closed loop, and the control reference signal that was used.

However, the comparison is still difficult because the studies use different reference signals. In the current work, we use RegA qualification and historical signals. Filtered ACE signals are used in [35], but this filtering is matched to the HVAC DR control system under test and not defined by the market or ISO/RTO. In [37], RegD reference signals are used in simulations.

Differences also exist between the experimental DR control systems and their representation in simulations. The simulations in [37] do not include any effects of communications latency. This omission likely has minimal effect on supply pressure/mass flow setpoint offset control in [37] because the number of control endpoints is small (a few AHUs) and the VAV dampers are not directly involved. The close comparison between the simulations of [37] and the experiments of [35] for "supply pressure/mass flow" in Table II provides some evidence for this conclusion. However, the high scores in [37] for the thermostat setpoint offset control (listed under "T-stat setpoint" in Table II) are questionable. In a large building, these control signals are sent to a large number of thermostat endpoints potentially creating significant latency. These effects lead to a lower delay score for the current work versus [37], 0.73 versus 0.99. As was argued earlier, this delay leads to lower precision and correlation scores and likely accounts for some of the difference between these scores for the current work and in [37], 0.38 versus 0.54 and 0.60 versus 0.92, respectively.

There are also significant differences between the experimental testbeds. The building that forms the testbed in [35] includes three AHUs, but the DR control is implemented on a single motor-fan unit that serves a single large auditorium through a single VAV. Also, the tests were carried out over a relatively short time period. Taken together, these two observations imply the HVAC system was not subject to significant exogenous disturbances that could degrade the DR control performance. In contrast, the current work is carried out on the four main AHUs serving nearly all the floor space of a 30 000-m^2 office building through nearly 500 VAVs. In addition, the experiments were carried out over 12 h, from a few hours after the morning start up to the beginning of the night setback. Because of this testbed configuration and test protocol, our system is subject to numerous exogenous perturbations, primary among which are highly variable occupancy, opening and closing of external air dampers for economizer operation, and the typical diurnal heat load cycle. These real-world processes likely act to decrease the correlation and precision performance scores for the current work in Table II, but illuminate the practical difficulties involved in qualifying a commercial HVAC system for DR participation.

In spite of the differences in the testbeds, test protocols and simulation fidelity, some general conclusions can still be drawn from the results in Table II. First, even if the BAS communication protocol is not particularly fast, the smaller number of control endpoints for the direct fan speed control or the supply pressure/mass flow control minimizes the effect of communication latency. In addition, by not involving VAV dampers, these control methods avoid this extra mechanical latency. These design choices are the likely reasons for the higher delay scores for the experiments in [35] as compared to the current work.

Closed-loop DR control is expected to perform better than open-loop control, however, the differences in Table II are not very large. The DR control in [37] is open loop, and although only studied in simulation, it shows tracking performance similar to the closed-loop control in [35]. The algorithm in [37] is based on a physics model of the duct system and fans. The quality of the tracking is a function of the accuracy of this model, which can be made quite detailed to capture the important effects. A similar conclusion can be made regarding the correlation score of the open-loop fan speed control in [36] versus [35]. However, as the open-loop control actuation moves closer to the endpoints, in particular the thermostats and VAVs in the thermostat setpoint offset controls, control performance suffers because accurate models are no longer feasible. The current work relies on a controller that utilizes reduced-order, data-driven statistical models developed in [34] which, by their nature, cannot capture all of the system detail. The reduced model accuracy may be reflected in Table II by the lower correlation and precision scores of the current work versus [35] and [37].

Throughout this analysis, it must be emphasized that the primary goal of the HVAC system is to maintain occupant comfort. The DR control system should allow the HVAC system to track, on average, its baseline power consumption curve over the day. Each of the DR controllers discussed above achieves this in slightly different ways. The closed-loop controls in [35] apply a bandpass filter to the incoming frequency regulation signal to eliminate the zero and low-frequency components. The existing HVAC control is then still able to track the low-frequency daily evolution of load. However, long-term testing was not carried out in the experiments of [35]. The open-loop supply pressure control in [37] does not directly control the VAVs. On longer time scales, the VAVs would be free to compensate for the changes in supply pressure, but the simulations in [37] were only carried out for 1 h and the nominal heat load on the HVAC system was constant over this time. The open-loop thermostat setpoint control method used in our testbed never allows the thermostat setpoint adjustments to deviate more than ± 2 °F. This approach naturally allows the HVAC system to track its baseline load over the entire day and was demonstrated over entire days in [34] and in the present work.

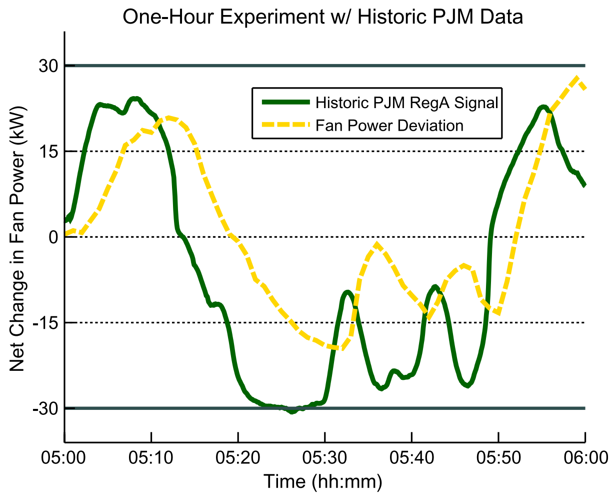

Fig. 6. ***An experiment using input from a 1-h sample of the PJM RegA signal recorded on May 5, 2014 (solid green line) and the response of the HVAC system fan power consumption (dashed yellow line). The less aggressive ramp rates of actual RegA market data compared to the RegA qualification test results in better tracking performance in the testbed building.***

D. Historical RegA Signal Tests

To assess performance under typical operating conditions, we also tested the same building and open-loop thermostat setpoint offset control method using historical RegA signals from the PJM demand response regulation market. The test protocol is the same as for the RegA qualification test except the duration of the historical RegA signal is 60 min instead of 40 min. A sample test is shown in Fig. 6. The distributions of delay, correlation, precision and aggregate performance scores are presented in Fig. 7. Compared to the RegA qualification test, the historical RegA signals display lower ramp rates. The delay scores are primarily a function of communication latency and are mostly unaffected by the lower

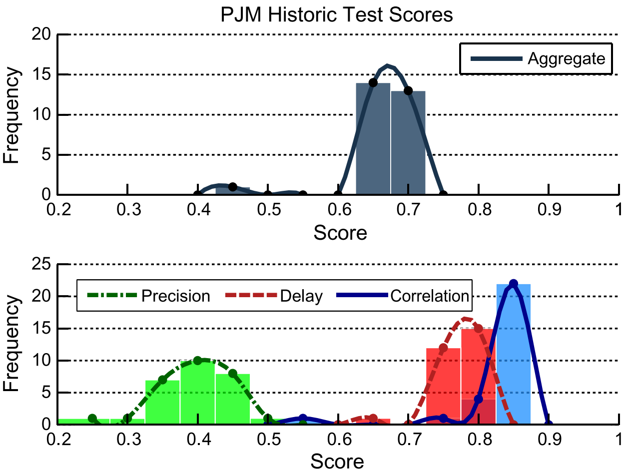

Fig. 7. ***Histograms of the PJM performance scores for the tests using historical RegA signals from the PJM market. Upper plot: Aggregate scores. Lower plot: Delay (red), Correlation (blue), and Precision (green) scores.***

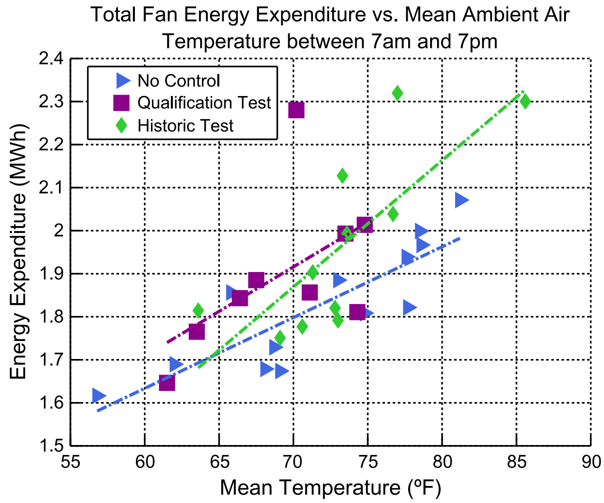

Fig. 8. ***Total fan energy consumption versus average daily temperature, both computed between 07:00:00 and 19:00:00. The data are sorted into days with no DR control, qualification test signal days, and historic signal days. Days where DR control was performed display higher energy consumption than when no DR control was performed.***

ramp rates. However, the lower ramp rates result in smaller control errors during the ramps and improvements in both the correlation and precision scores. The aggregate score is improved to $\sim$0.65 which is better than the minimum aggregate score of 0.40 required for continued participation in the PJM market.

VI. COST OF CONTROL: ENERGY COSTS

DR may have several costs, including capital costs to install the controls and communications, and operating cost arising from maintenance. There may, however, also be operating costs related to increased energy consumption due to less efficient load operation. Estimates of the capital costs are beyond the scope of the current work. Instead, we focus on the question of operational cost: If the capital cost can be sufficiently reduced, how will the profitability of DR be affected by reduced load efficiency? Specifically, we consider the impact of DR on the time-average efficiency of commercial HVAC operations and the cost of the additional energy required to operate the HVAC system while being controlled for DR. We note that the impact of infrequent DR control, such as time-of-use tariffs and peak shaving, on time-average load efficiency and energy costs are expected to be small compared to the high value of these operations [4], [46]. However, the frequency regulation considered here is expected to operate frequently and even relatively minor changes in efficiency could significantly increase energy costs compared to the expected frequency regulation revenue.

The importance of the impact of frequency regulation on HVAC efficiency can be understood by considering a traditional generator providing the same service. A traditional generator dispatched to a constant power output P_0 will consume fuel at a constant heat rate H_0. If that same generator is also used for frequency regulation with capacity ΔP, its output P will continually vary between $P_0 - \Delta P$ and $P_0 + \Delta P$, but its time-average power output will still be the same, i.e., $\overline{P} = P_0$ (assuming a balanced regulation signal). However, the continual ramping of the generator reduces its efficiency and increases its time-average heat rate, i.e., $\overline{H} > H_0$. The higher heat rate increases the generator operating cost, and this increase is factored into the generator's bid [16] into the frequency regulation market.

Here, we perform experiments more like a traditional generator providing frequency regulation. We compute the total fan energy consumption during the 07:00:00 to 19:00:00 time window for days when the PJM qualification test (described in Section V-B) and historical signals (Section V-D) are applied and during days when no testing is performed. We only consider summer, nonholiday weekdays to compare days with similar load profiles. Fig. 8 displays each day's total fan energy versus the average daily temperature (also computed between 07:00:00 and 19:00:00). The figure shows a linear fit to the data for each type of reference signal.

Given that each datapoint represents an entire day of testing, the total number of datapoints is limited, and there is considerable scatter in the data. Despite this, Fig. 8 displays a general trend in which the energy consumption for qualification and historical test days is higher than when no DR control is performed. At an average ambient temperature of 75 °F, the extra energy is ∼100 kW-h corresponding to \$10 for a tariff of \$0.10/kW-h. During the tests with the historical RegA signal, the building provided 30 kW of frequency regulation capacity for 6 h. Based only on the cost of the extra energy, the operating cost to provide this frequency regulation service is $(1000/(30 \times 6)) \times 10 =\sim \$55$ /MW/h, which is higher than typical PJM clearing prices [47], underscoring the significant impact that energy inefficiencies could have on frequency regulation profitability.

While the results displayed in Fig. 8 are interesting, there is not enough data present to draw firm conclusions on the energy efficiency of commercial HVAC frequency regulation. A more rigorous analysis of the inherent energy costs is presented in [48], which utilizes the same testbed building described in the current work. An analogy is drawn between the HVAC energy penalty and a lossy battery providing "energy neutral" frequency regulation. In this case, the internal losses during the balanced charge/discharge cycles dissipate the battery's stored energy causing the battery's average state of charge (SOC) to continually decrease. Additional charging is needed to maintain the SOC creating a source of operating cost. In [48], a series of experiments was performed to capture this effect for our commercial HVAC testbed—a series of step-like thermostat setpoint offset commands were used to create a nearly energy neutral interaction between the HVAC fans and the grid. After releasing the control, the HVAC had to consume additional energy to return to the nominal baseline operation. In analogy with batteries, the electrical round trip efficiency of the HVAC "charging and discharging" of the thermal mass of the building was found to be ∼0.46 and was used to estimate the operating cost of the frequency regulation service at ∼\$30 /MW/h for an energy tariff of \$0.10 /kW-h.

We emphasize that both the results in Fig. 8 as well as those described in [48] are for a single commercial HVAC system using one type of DR control (open-loop thermostat setpoint offsets). Additional characterization work is needed to determine the general trends in the costs of additional energy consumption and to explore means for reducing this additional cost.

VII. CONCLUSION

Growth in renewable generation is challenging operational strategies that have traditionally maintained the balance between generation and load. Flexible demand-side resources may help alleviate this issue. Ancillary services markets such as PJM's frequency regulation market are expanding to allow these resources to participate, provided they can pass minimum performance tests. This new application of DR requires controlling these resources on faster time scales, far more frequently, and more accurately than is the case in traditional DR applications.

We have reviewed experimental and simulation studies [34]–[37] and presented our own experimental work investigating the ability of a commercial HVAC system to provide frequency regulation services. We have used these studies to compare different DR control methodologies (open versus closed loop) and different DR control inputs (direct fan speed offset, supply pressure/mass flow offset, and thermostat setpoint offset). A small number of studies reported to date have been undertaken under very different conditions, including bench top experimental tests of single components in [36], in-building experimental test and simulation of small sections of a HVAC system in [35], numerical simulation of representative portions of a building in [37], and full-scale experimental tests on an entire building in [34], and in the experiments presented here. Table II provides a summary of these studies.

Although these projects demonstrate wide diversity, some general conclusions can still be drawn. First, in-building communication latency and mechanical latency can significantly impact the performance of fast DR controls. Comparing the current experimental results with those in [37] shows that if these latencies are not considered, DR control performance may be significantly overestimated. Second, the choice of control inputs can impact the degree to which these latencies affect performance, from minimal effect for direct fan speed control [35], [36] or supply pressure/mass flow setpoint offsets [35], [37] to a significant effect for thermostat setpoint offsets in [34] and the present work. However, the choices that offer reduced latency come at a cost of not being able to customize the DR participation for individual conditioned spaces or occupants. Third, as expected, closed-loop DR controls [35] generally perform better than the open-loop DR controls in [34], [36], [37] and the present work. However, as the control input is moved closer to the HVAC fan (e.g. direct fan control in [36] or supply pressure/mass flow setpoint offsets in [37]), open-loop performance appears to approach that of closed-loop control. We note again, though, that this improvement in performance comes at the cost of not being able to customize the DR participation for individual conditioned spaces or occupants, as can be achieved with the methods presented in the current experimental studies.

The new experiments presented in this paper and the related experiments in [48] also point to an issue that may have significant impact on the economic viability of fast, continuously operated DR control like the frequency regulation service explored here. A commercial HVAC system that is continually perturbed from nominal operation, and that is rapidly ramped between very different operating states, will display a lower average efficiency than when operated in a steady manner. The lower efficiency translates into increased average energy consumption. The effect is similar to that experienced by a traditional fossil-fueled generator whose efficiency suffers when ramped up and down to provide frequency regulation service. Much like a traditional generator, the increased cost of serving the HVAC load should be incorporated into the cost of providing the frequency regulation service. For the building and control system investigated here, we find this cost rather high (∼$55 /MW/h). In contrast, for relatively infrequent DR control like spinning reserves or peak shaving, the impact of these effects on long-run efficiency and DR economics is minimal.

There are many interesting and important directions for future investigations focused on using commercial HVAC systems for fast, continuously operated DR control like frequency regulation. Although we have been able to make some general comparisons, the diversity of test platforms and protocols precludes detailed comparisons of control methods and control inputs. Current and future research would greatly benefit from access to several standard test platforms comprising real commercial buildings spanning typical commercial HVAC architectures. Furthermore, standardization of test protocols and performance metrics, e.g., a representative sample of PJM historical RegA and RegD signals combined with PJM performance metrics [16], would enable more meaningful comparison of different control methods. A third area that deserves attention is understanding and minimizing the impact that persistently excited DR control, such as frequency regulation, has on HVAC energy consumption and the economics of these forms of DR. These impacts should also be explored for other resources that are being considered for such ancillary services, e.g., commercial refrigeration, residential HVAC, and lighting. ■

Acknowledgment

The authors would like to thank P. Barooah and S. Kiliccote for useful discussions.

REFERENCES

[1] S. Borenstein, "The long-run efficiency of real-time electricity pricing," *Energy J.*, vol. 26, no. 3, pp. 93–116, 2005.

[2] R. Sioshansi and W. Short, "Evaluating the impacts of real-time pricing on the usage of wind generation," *IEEE Trans. Power Syst.*, vol. 24, no. 2, pp. 516–524, 2009.

[3] H. Allcott, "Rethinking real-time electricity pricing," *Resource Energy Econ.*, vol. 33, no. 4, pp. 820–842, 2011.

[4] D. J. Hammerstrom *et al.*, "Pacific Northwest GridWise Testbed Demonstration Projects. Part I. Olympic Peninsula Project," Pacific Northwest National Laboratory, Tech. Rep. PNNL-17167, 2007.

[5] M. Roozbehani, M. A. Dahleh, and S. K. Mitter, "Volatility of power grids under real-time pricing," *IEEE Trans. Power Syst.*, vol. 27, no. 4, pp. 1926–1940, 2012.

[6] C. W. Taylor, "Concepts of undervoltage load shedding for voltage stability," *IEEE Trans. Power Delivery*, vol. 7, no. 2, pp. 480–488, 1992.

[7] S. Arnborg, G. Andersson, D. J. Hill, and I. A. Hiskens, "On undervoltage load shedding in power systems," *Int. J. Electr. Power Energy Syst.*, vol. 19, no. 2, pp. 141–149, 1997.

[8] P. Palensky and D. Dietrich, "Demand side management: Demand response, intelligent energy systems, smart loads," *IEEE Trans. Ind. Inf.*, vol. 7, no. 3, pp. 381–388, 2011.

[9] N. Beute and G. Delport, "An historic overview of controlling domestic water heating," in *Proc. 14th Domestic Use Energy Conf.*, 2006, pp. 41–46.

[10] B. Hastings, "Ten years of operating experience with a remote controlled water heater load management system at Detroit Edison," *IEEE Trans. Power Apparatus Syst.*, vol. PAS-99, no. 4, pp. 1437–1441, 1980.

[11] W. Abercrombie, S. Bartlett, and B. Murray, "Load management developments in Queensland," in *Proc. Eng. Conf.*, Australia, 1980, pp. 370.

[12] O. Ma *et al.*, "Demand response for ancillary services," *IEEE Trans. Smart Grid*, vol. 4, pp. 1988–1995, Dec. 2013.

[13] M. Kintner-Meyer, "Regulatory policy and markets for energy storage in North America," *Proc. IEEE*, vol. 102, no. 7, pp. 1065–1072, Jul. 2014.

[14] 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," Tech. Rep. LBNL-5958E, Nov. 2012.

[15] J. Hurley, P. Peterson, and M. Whited, "Demand response as a power system resource," The Regulatory Assistance Project, Tech. Rep., May 2013.

[16] PJM Regional Transmission Operator, "PJM Manual 12: Balancing Operations," Forward Market Operations, Tech. Rep., Rev. 32, Apr. 2015.

[17] W. Bower *et al.*, "The advanced microgrid: Integration and interoperability", Sandia National Laboratories, SAND2014-1535, Mar. 2014.

[18] S. Chanda, V. Venkataramanan, and A. Srivastava, "Real time modeling and simulation of campus microgrid for voltage analysis," in *Proc. North Amer. Power Symp.*, Sep. 2014, DOI: 10.1109/NAPS.2014.6965481.

[19] PJM Regional Transmission Operator, "PJM Manual 11: Energy and Ancillary Services Market Operations," Forward Market Operations, Tech. Rep., Rev. 68, Aug. 2014.

[20] A. Wood, B. Wollenberg, and G. Sheblé, *Power Generation, Operation and Control*, 3rd ed. New York, NY, USA: Wiley, 2014.

[21] N. Jaleeli, L. VanSlyck, D. Ewart, L. Fink, and A. Hoffmann, "Understanding automatic generation control," *IEEE Trans. Power Syst.*, vol. 7, pp. 1106–1122, Aug. 1992.

[22] Supreme Court of the United States, "Case No. 18-841: Federal Energy Regulatory Commision vs. Electric Power Supply Association."

[23] Federal Energy Regulatory Commission (FERC), "FERC Order 745: Demand Response Compensation in Organized Wholesale Energy Markets," Tech. Rep., Mar. 2011.

[24] Federal Energy Regulatory Commission (FERC), "FERC Order 755: Frequency Regulation Compensation in the Organized Wholesale Power Markets," Tech. Rep., Oct. 2011.

[25] M. Klobasa, "Analysis of demand response and wind integration in Germany's electricity market," *IET Renew. Power Gen.*, vol. 4, no. 1, pp. 55–63, 2010.

[26] P. Cappers, C. Goldman, and D. Kathan, "Demand response in US electricity markets: Empirical evidence," *Energy*, vol. 35, no. 4, pp. 1526–1535, 2010.

[27] D. S. Callaway, "Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy," *Energy Conv. Manage.*, vol. 50, no. 5, pp. 1389–1400, 2009.

[28] S. Koch, J. L. Mathieu, and D. S. Callaway, "Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services," in *Proc. Power Syst. Comput. Conf.*, 2011, pp. 1–7.

[29] S. Kundu, N. Sinitsyn, S. Backhaus, and I. Hiskens, "Modeling and control of thermostatically controlled loads," in *Proc. Power Syst. Comput. Conf.*, 2011.

[30] N. Lu, D. P. Chassin, and S. E. Widergren, "Modeling uncertainties in aggregated thermostatically controlled loads using a state queueing model," *IEEE Trans. Power Syst.*, vol. 20, no. 2, pp. 725–733, 2005.

[31] S. Kundu, and I. A. Hiskens, "State-space modelling of hysteresis-based control schemes," in *Proc. Eur. Control Conf.*, 2013, pp. 2535–2540.

[32] M. A. Piette *et al.*, "Development and evaluation of fully automated demand response in large facilities," Lawrence Berkeley National Laboratory, Tech. Rep. CEC-500-2005-013, 2005.

[33] N. Motegi, M. A. Piette, D. S. Watson, S. Kiliccote, and P. Xu, "Introduction to commercial building control strategies and techniques for demand response," Lawrence Berkeley National Laboratory, Tech. Rep. LBNL-59975, 2007.

[34] G. Goddard, J. Klose, and S. Backhaus, "Model development and identification for fast demand response in commercial HVAC systems," *IEEE Trans. Smart Grid*, vol. 5, no. 4, pp. 2084–2092, 2014.
[35] Y. Lin, P. Barooah, S. Meyn, and T. Middelkoop, "Experimental evaluation of frequency regulation from commercial building HVAC systems," *IEEE Trans. Smart Grid*, vol. 6, pp. 776–783, Mar. 2015.
[36] J. MacDonald and S. Kiliccote, "Commercial building loads providing ancillary services in PJM," presented at the *ACEEE Summer Study Energy Efficiency Buildings*, Pacific Grove, CA, USA, Aug. 17–22, 2014.
[37] P. Zhao, G. P. Henze, S. Plamp, and V. J. Cushing, "Evaluation of commercial building HVAC systems as frequency regulation providers," *Energy Buildings*, vol. 67, pp. 225–235, 2013.
[38] J. Murphy, "High-performance VAV systems," *ASHRAE J.*, vol. 53.10, pp. 18–28, Oct. 2011.
[39] P. Barooah, A. BuŠic, and S. Meyn, "Spectral decomposition of demand-side flexibility for reliable ancillary services in a smart grid," in *Proc. 48th Annu. Hawaii Conf. Syst. Sci.*, 2015, pp. 2700–2709.
[40] J. E. Braun, K. W. Montgomery, and N. Chaturvedi, "Evaluating the performance of building thermal mass control strategies," *HVAC&R Res.*, vol. 7, no. 4, pp. 403–428, 2001.
[41] J. C. Lam, D. H. Li, and S. Cheung, "An analysis of electricity end-use in air-conditioned office buildings in Hong Kong," *Building Environ.*, vol. 38, no. 3, pp. 493–498, 2003.
[42] B.-B. Lu, X.-Y. Chen, X.-H. Ding, Y.-C. Liao, and K. Yu, "A direct control of central air conditioning load considering indoor comfort," *Environ. Energy Sustain. Develop.*, vol. 4, pp. 369, 2013.
[43] A. Mammoli, H. Barsun, R. Burnett, J. Hawkins, and J. Simmins, "Using high-speed demand response of building HVAC systems to smooth cloud-driven intermittency of distributed solar photovoltaic generation," in *Proc. IEEE PES Transm. Distrib. Conf. Expo.*, 2012, DOI: 10.1109/TDC.2012.6281419.
[44] G. S. Ledva, E. Vrettos, S. Mastellone, G. Andersson, and J. L. Mathieu, "Applying networked estimation and control algorithms to address communication bandwidth limitations and latencies in demand response," in *Proc. 48th Hawaii Int. Conf. Syst. Sci.*, 2015, pp. 2645–2654.
[45] L. Ljung, *System Identification. Theory for the User*, 2nd ed. Englewood Cliffs, NJ, USA: Prentice-Hall, 1999.
[46] W. J. Cole, K. M. Powell, E. T. Hale, and T. F. Edgar, "Reduced-order residential home modeling for model predictive control," *Energy Buildings*, vol. 74, pp. 69–77, 2014.
[47] PJM, "PJM regulation zone preliminary billing data," Apr. 2013. [Online]. Available: http://pjm.com/markets-and-operations/market-settlements/preliminary-billing-reports/pjm-reg-data.aspx
[48] I. Beil, I. Hiskens, and S. Backhaus, "Round-trip efficiency of fast demand response in a large commercial air conditioner," *Energy Buildings*, vol. 97, pp. 47–55, 2015.

ABOUT THE AUTHORS

Ian Beil received the B.S. degree from Washington University in St. Louis, St. Louis, MO, USA, in 2010, and the M.S. and Ph.D. degrees from the University of Michigan, Ann Arbor, MI, USA, in 2013 and 2015, respectively, all in electrical engineering.

His research interests include power systems topics such as demand response, load flow analysis, and grid integration of renewable generation and electric vehicles. He is currently an engineer in the Electrical Analytical Division at Sargent & Lundy, LLC, Chicago, IL, USA.

Ian Hiskens (Fellow, IEEE) is the Vennema Professor of Engineering in the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI, USA. He has held prior appointments in the Queensland electricity supply industry, and various universities in Australia and the United States. His research interests lie at the intersection of power system analysis and systems theory, with recent activity focused largely on integration of renewable generation and controllable loads.

Dr. Hiskens is actively involved in various IEEE societies, and is VP-Finance of the IEEE Systems Council. He is a Fellow of Engineers Australia and a Chartered Professional Engineer in Australia.

Scott Backhaus received the Ph.D. degree in physics from the University of California at Berkeley, Berkeley, CA, USA, in 1997.

He came to Los Alamos National Laboratory, Los Alamos, NM, USA, in 1998 and was the Director's Funded Postdoctoral from 1998 to 2000, a Reines Postdoctoral Fellow from 2001 to 2003, and a Technical Staff Member from 2003 to the present. Recently, his research interests have focused on energy-related topics, including the fundamental science of geologic carbon sequestration and grid-integration of renewable generation.

Dr. Backhaus received an R&D 100 Award in 1999 and the Technology Reviews Top 100 Innovators Under 35 Award in 2003.