



Round-trip efficiency of fast demand response in a large commercial air conditioner



I. Beil^{a,*}, I.A. Hiskens^a, S. Backhaus^b

^a Department of Electrical Engineering and Computer Science, University of Michigan, 1301 Beal Ave., Ann Arbor, MI 48109-2122, United States

^b Materials Physics and Applications Division, Los Alamos National Laboratory, Los Alamos, NM 87545, United States

ARTICLE INFO

Article history:

Received 23 April 2014

Received in revised form 10 February 2015

Accepted 14 March 2015

Available online 2 April 2015

Keywords:

Demand response

Ancillary services

Energy efficiency

Heating ventilation and air conditioning

ABSTRACT

Air conditioning (AC) of large commercial buildings represents an attractive target for many different forms of demand response (DR) including DR for ancillary services (AS) such as frequency regulation. The operating cost of such DR is typically discussed in terms of occupant discomfort. However at fast timescales, perturbations to well-functioning building controls may increase the total energy consumption relative to a baseline that does not provide DR ancillary services (DR-AS). The extra energy is a cost of control to the asset owner and should be factored into the cost of providing DR-AS. We performed DR experiments on a ~30,000 m² office building, and at the 15-min time scale of these experiments, we find the extra energy consumption to be significant. Similar to battery energy storage, we express the energy cost in terms of a round-trip efficiency and use this metric in a simple economic analysis of the cost of frequency regulation from these resources and discuss potential impacts on advanced load management methods.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

The penetration of renewable generation is increasing significantly in many electrical grids. In the U.S., much of this expansion is being driven by renewable portfolio standards (RPS) [1] that mandate a certain fraction of generation be derived from renewable resources, with wind and photovoltaic (PV) generation being the primary contributors. Both of these resources are variable and stochastic, i.e. they have limited controllability. As the penetration of renewable generation increases, there is less room in the generation stack for controllable generation for mitigating renewable fluctuations. Viewed in this way, even the expansion of flexible natural gas generation in many parts of the U.S. may ultimately be limited in its ability to provide these balancing resources. Other controllable resources will be required to provide such ancillary services (AS).

Grid-level energy storage is one candidate to provide these services. Pumped hydro storage is an economical, highly efficient form of energy storage, however, its availability is limited by local geography [2]. Some forms of battery energy storage can have roughly similar efficiency, however, the capital costs remains stubbornly

high. The lack of universally available storage has encouraged the development of demand response (DR). The intent of DR load control is to adjust the time when energy is consumed to assist with the management of generation-load balance. DR is used in many electrical grids to provide demand response ancillary services (DR-AS) with the majority of these uses being spinning reserves [3], i.e. DR-AS that is called on relatively infrequently. The impact to the interrupted load is primarily disruption of the end use. Other forms of DR-AS, such as frequency regulation service [4], will require much more frequent and continual modification of load behavior which may have performance implications beyond just a disruption of the end use.

The assumption by many DR research efforts [4–7] is that disruption of the load by DR has little or no effect on the total energy consumption or that the effect is insignificant compared to the value of the DR service. For loads that are simply deferred, for example dish washing and clothes drying, or for grid services that only require a few load adjustments per day, this assumption is likely quite good. However, if DR-AS requires that load consumption be modulated about a mean level on sub-hourly timescales, this assumption may not hold and additional energy consumption may be required to serve the disturbed load. Some experimental [8] and simulation [9] work has demonstrated an increase in average total energy consumption due to the relatively slow preheating or precooling of homes driven by time variable energy pricing. However, insufficient attention has been given to possible

* Corresponding author. Tel.: +1 734 646 2606.

E-mail addresses: ianbeil@umich.edu (I. Beil), hiskens@umich.edu (I.A. Hiskens), backhaus@lanl.gov (S. Backhaus).

energy losses incurred because of DR control and especially DR-AS control.

Large commercial air conditioning (AC) systems are attractive loads for DR and DR-AS applications because they already include relatively sophisticated control and communications architectures, reducing the expected incremental capital cost of adding DR capability. In contrast, a population of small residential loads may provide a well-understood and repeatable demand profile [10], but the capital cost of the enabling control and communication infrastructure is significant. The operational or control cost of AC-based DR is typically discussed in terms of occupant discomfort – a very important metric because a DR control that exceeds comfort levels too often could be discontinued. Occupant discomfort depends upon the DR control system and how it weights discomfort versus other objectives, e.g. accuracy in tracking load power reference signals. However, there may exist additional operating costs. If the building controls for a large commercial AC system are operating in a quasi-steady state, perturbing these operations will likely increase the time-averaged energy consumption, thereby creating additional operating costs for the asset owner who is providing the DR-AS service. This manuscript focuses on characterizing these additional energy costs for one type of large commercial building AC at timescales comparable to those in frequency regulation service.

We have performed DR-AS experiments on a ~30,000 m² office building to develop a methodology for measuring the extra energy consumption, to determine this extra consumption for one class of commercial buildings using one type of AC system, and to develop an economic model for the variable operating cost of the DR-AS. The DR experiments are primarily carried out at the 15-min time scale, a speed appropriate for mitigating fluctuations due to wind and PV variations, and approaching the control speeds needed for participation in frequency regulation markets [11]. We characterize this AC-based DR in terms of a round-trip efficiency so that it can be directly compared to other forms of energy storage, e.g. the round-trip energy losses incurred when charging and discharging a battery.

The remainder of this manuscript is organized as follows. Section 2 presents a review of relevant DR literature. Section 3 provides a brief description of the AC system used for experimentation, sufficient for understanding the main results. Section 4 introduces an analogy to round-trip efficiency of battery charging/discharging that will be useful in the discussion of our experiments. Section 5 describes the experimental protocol used for determining AC-based DR round-trip efficiency – a method we believe is applicable to most large commercial AC systems and large aggregates of residential AC units. Section 6 presents the results of our measurements and a comparison with additional analysis of data from related experiments [8] and simulations [9]. Section 7 provides a discussion of our results in terms of an economic model for the DR service and some potential implications for advanced load control. Finally, Section 8 presents our conclusions and a discussion of potential future work.

2. Literature review

The physical structure of large commercial AC systems is interconnected such that control actions taken by one component affect the behavior of all attached components, often leading to complex dynamics. For control purposes, many researchers describe the behavior of these systems using reduced-order models, e.g. Ziegler-Nichols tuning [12] or self-learning neural networks [13]. Such approaches allow the operator to aggregate the many individual control points and devices in the system into a few key control setpoints. In previous work [14], we take a related approach that represents the state of a large number of similar devices

(variable air volume units, described in Section 3) using a probability distribution, and relates this distribution to the AC fan power consumption. We also developed a model of the dynamics of this distribution and used this predictive model in an open-loop method capable of providing DR-AS at speeds approaching those required for participation in frequency regulation markets. We use this same open-loop control strategy [14] in the current work.

The work in this manuscript also builds on previous experimental work on DR in commercial AC systems. This includes [15] which demonstrated load shedding by AC systems across geographically separated commercial buildings when exposed to increases in a dynamic electrical price on hour-long timescales. In a separate test presented in [16], the results in [15] were validated on a larger population of buildings subject to critical peak electricity pricing. Experiments using customer AC loads to provide spinning reserve services [17] have also been conducted. The Open Automated Demand Response (OpenADR) communication protocol [18] has been used to control the AC load of buildings [19] by sending both price and load curtailment signals. However, this previous work is primarily focused on longer time scales and infrequently called DR services. Our focus in this manuscript is on faster and continuously controlled DR-AS such as frequency regulation.

Some recent work has focused on more advanced AC control paradigms including genetic algorithms [20], occupancy-based control of entire building floors [21], model predictive control (MPC) to reduce transient loads [22], and MPC in conjunction with weather forecasting to regulate an entire building AC system [23]. Predictive control schemes have also been used to minimize building energy costs under time-of-use pricing [9,24] as well as under dynamic pricing scenarios [9,25]. Larger system studies have investigated the impact of DR on entire electrical distribution circuits [26]. However, most of this work is focused on control at slow time scales rather than the fast frequency regulation we consider here. Several efforts used simulations [4,27] and experiments [27,28] to examine commercial AC response to frequency regulation signals on shorter-term timescales, but these studies did not include the loss mechanisms examined in our experiments.

3. AC system description

The details of commercial AC systems vary from building to building, with the larger systems often being custom designed to the particular application. However, many large systems share a common underlying structure. The AC system in this work (see Fig. 1) is comprised of a central chiller plant that distributes chilled water to the heat exchangers in several independent air handling units (AHU). The fans in the AHUs circulate building air through the heat exchangers and to the conditioned spaces. The flow of cold air into each space is regulated by individual variable air volume (VAV) units. The rest of this section provides details of the AC system that are important for understanding the properties of the DR-AS experiments. For a more in depth discussion of the HVAC layout, see our previous work [14].

For the purposes of this discussion, we begin at the VAVs which are the air inlets to the conditioned spaces (see Fig. 1). A damper valve in the VAV adjusts to regulate the amount of cooled air flowing from the supply air duct into a conditioned space. Unlike the discrete, hysteretic control of small AC units, the error signal from the conditioned-space thermostats is used by a local proportional-integral-derivative (PID) controller to continuously vary the VAV damper valve opening. The damper valves can be actuated between 100% (fully open) down to about 20–30% [29]. The lower limit is set to ensure the conditioned space always receives a required minimum level of ventilation. The local PID controller and the mechanical response time of the damper valve opening determine

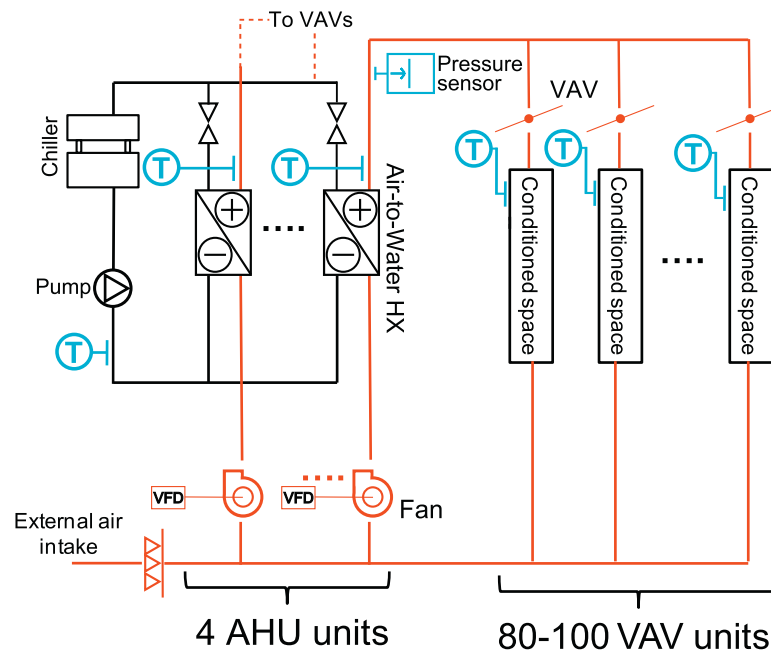


Fig. 1. A schematic drawing of the AC system used in this work. The AC system is composed of two main sets of loops. The first is a water loop that circulates water between the chiller plant and air-to-water heat exchangers where the water loop interfaces with the air loops. In the air loops, fans in the Air Handling Units (AHU) force warm air returned from the conditioned spaces through the air-to-water heat exchangers where it is cooled and then delivered to the inlets of an array of variable air volume (VAV) units. Thermostats (T) in the conditioned spaces provide input to a controller in the VAV to regulate the amount of cooled air that enters each 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 the pressure of the cooled air supplied to the VAVs.

how quickly the air flow responds to changes in the temperature error signal from the thermostat. This is in addition to any communication latency. The VAVs in the AC system in this work typically respond to large changes in conditions in ~ 1 min.

VAVs that are physically close to each other are grouped together and connected via ductwork back to a common AHU (see Fig. 1). In the AC system used in this work, there are four AHUs each serving groups of 80–100 VAVs. The AHUs provide two vital functions. First, the air-to-water heat exchanger in each AHU removes heat from the warm return air so the air supplied to the VAVs is cool. Second, a local control in each AHU provides a steady supply air pressure to the VAVs. If the thermostats in a group of VAVs are calling for more cooling, the local PID loops open the VAV damper valves to increase the air flow. Without additional controls, the air pressure on the downstream side of the supply air fan (and at the VAV inlets) would fall, counteracting the desired changes. However, a signal from a pressure sensor at the fan output is used in a local control loop that regulates the supply air fan speed via a variable frequency drive (VFD) to maintain a constant outlet pressure. This control action results in an increase in electrical fan power when the thermostats call for more cooling. A decrease in fan power occurs when the thermostats call for less cooling. The control loop on the fan VFDs responds on approximately the same time scale of the VAVs, i.e. in ~ 1 min. The total nameplate capacity of the four AHUs in the LANL testbed is 295 kW.

A chilled water loop removes heat from the warm return air and the warmed water is returned to a chilled water plant that removes the heat from the water and finally rejects it to ambient. There are two types of control on the chilled water loop. First, a water valve at the inlet to each water-to-air heat exchanger is automatically controlled to fix the AHU outlet supply air temperature. A second control loop fixes the water outlet temperature from the chilled water plant by adjusting the number of chiller compressor units that are engaged. Both of these controls operate on a time scale of ~ 10 min making them relatively slow compared to the VAV and supply air pressure controls.

The Building Automation System (BAS) provides communication and control for the entire HVAC system. Based on data gathered from a building's sensors and actuators, the BAS determines the nominal operating set-points for supply air duct static pressure, AHU supply air temperature, and the chilled water supply temperature. The BAS updates these set-points on a timescale slow compared to the DR-AS control.

4. Battery analogy

When discussing the energy efficiency of our AC-based DR ancillary service experiments, it is convenient to make an analogy to the losses in a battery and the round-trip efficiency of battery charging/discharging. In Fig. 2, we plot the grid power exchange and state of charge (SOC) for a hypothetical charge–discharge cycle for a lossy battery. In control period 1, the battery is charged with a constant power P drawn from the grid. Because of charging losses, the power delivered to the battery is less than P . The battery SOC increases but at a rate lower than P . During the control period 2 (the same duration as period 1), the battery is discharged with a constant power P delivered to the grid. Because of discharge losses, the battery discharges at a rate higher than P , and the battery SOC decreases at a rate faster than P .

At the end of control period 2, the net energy exchange with the grid is zero, however, the internal state of the battery (i.e. the SOC) is lower than at the beginning of this process. To restore the internal state requires an additional charging period (control period 3) to replace the energy dissipated in the first two periods. After restoring the SOC to its initial value, the battery has taken more energy from the grid (E_{in}) than it has delivered to the grid (E_{out}). The round-trip efficiency for this battery can now be defined as

$$\eta_{RT} = \frac{E_{out}}{E_{in}}. \quad (1)$$

The same definition works equally well for more complicated charge/discharge processes, e.g. the charge and discharge periods

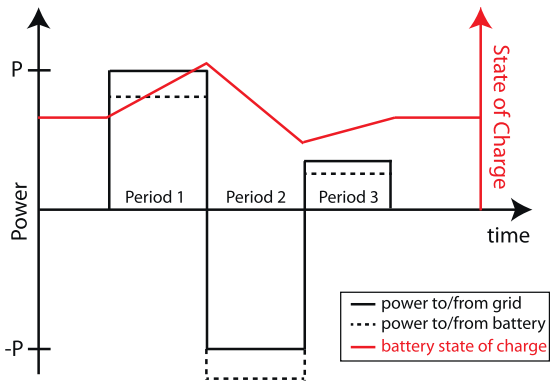


Fig. 2. Prototypical charge/discharge cycle of a lossy battery. During the first period, the battery charges by drawing power P from the grid (solid black line), however, the losses reduce the power delivered to the battery (dashed black line) and the battery state of charge (SOC) (red line) increases at a reduced rate. During the second control period, the battery discharges with power P delivered to the grid. Losses result in the battery discharging at a rate larger than P and the SOC decreases at an accelerated rate. This charge/discharge cycle results in the SOC being lower than at the start of the process. A third period of charging is required to restore the SOC to the starting value. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

need not be the same length or exchange the same power with the grid. The only requirement is that the battery SOC is the same at the beginning and end of the process, and E_{in} and E_{out} account for all of the energy extracted from or delivered to the grid, respectively.

5. Experimental protocol

To control the building power consumption in this work, we leverage an open-loop DR control algorithm developed for the same building in previous work [14]. The algorithm converts an input electrical power (kW) change command into a Global Thermostat Reset (GTR) that simultaneously shifts nearly all of the ~ 500 thermostats in the building to approximate the desired electrical power change from the fans in the AHUs (see Fig. 1). The extreme reduction in dimensionality of the control model limits the accuracy of

the algorithm [14], but it is sufficient for the experiments carried out in this work.

The BAS is programmed to perform a repetitive set of DR-AS experiments throughout the day. A single-day sample of the AC fan power recordings is presented in Fig. 3. The Control Input trace indicates the DR-AS signal implemented by the BAS. It consists of a 15-min period where the AC fan power is increased by 15 kW above the baseline followed by a 15-min period where the AC fan power is decreased by 15 kW below the baseline. This near zero-energy exchange with the grid is followed by a 90-min period where the DR-AS control is set to zero to allow the building to return to its baseline operation before the next DR-AS experiment is initiated.

These experiments are repeated throughout the day for approximately 15 days. To check for hysteresis effects, the experiment is run with both polarities of the DR-AS control pulse. The amplitude of the DR-AS control pulse is 30 kW in the majority of these experiments. To determine an accurate baseline, the AC fan power data is interpolated between a short time window that immediately precedes the DR-AS control pulse and a second short time window immediately prior to the next DR-AS control pulse. Subtracting the baseline from the recorded fan power yields the DR power provided to the grid (see Fig. 3).

To assess the effect of baseline accuracy on our results for round trip efficiency, this same linear interpolation was performed over two hour time periods on seven days in which control actions were not taken, i.e. the AC system was unperturbed. The tests revealed an average energy over the baseline of 0.724 kW h per two hour control period, and an average energy under the baseline of 0.532 kW h for the same time frame. Our experiments use control pulses of $30 \text{ kW} \cdot 0.25 \text{ h} = 7.5 \text{ kW h}$ giving a conservative estimate of the error in round trip efficiency on the order of $\sim 10\%$. (For a more in-depth examination of variability in baseline AC load estimation, see e.g. [30]).

Fig. 4 shows a magnified view of the AC fan power over a single control period (from a different experiment than shown in Fig. 3). The experiment is initiated at 11:00 using a 30 kW amplitude DR control signal with the same two 15-min control periods as discussed above. During the first control period, the extra power consumption is converted into extra cooling power in excess of the internal heat generation and the heat leak into the building.

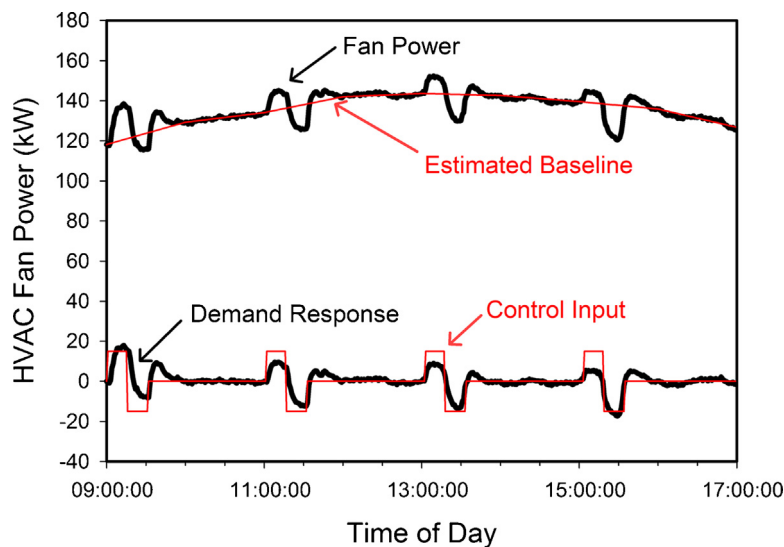


Fig. 3. Full-day recording of the DR control input signal, the AC fan power data, linearly interpolated baseline load, and resulting demand response (DR) profile. Each control period consists of a 15-min interval where an open-loop control [14] is used to increase the AC fan power above the baseline by 15 kW followed by a second 15-min period where the AC fan power is decreased below the baseline by 15 kW. Following this near zero-energy exchange with the grid, the DR control is set to zero for 90 min to allow the building to return to its baseline operation before the next DR experiment is initiated. DR experiments are not performed between 5 pm and 9 am because night setback and morning cool down of the building take precedence.




Fig. 4. A typical DR control transient. The black dashed line is the DR control signal implemented by the Building Automation System (BAS) which translates it to a global thermostat shift as in [14]. The control signal is zero for 90 min prior to 11:00 (not shown), i.e. the control is requesting no change relative to the fan power baseline and the building AC is running as though no demand response is requested. After the positive and negative cycles of DR, the control is released, i.e. no DR is requested and the building is allowed to return to its baseline operation. The solid blue trace is the measured fan power response. The fan power tracks the DR control signal reasonably well with a few minutes of time delay. After 11:30, the measured fan power remains high, i.e. above its baseline power, to return the building to its state prior to the initiation of the DR control at 11:00 h. The dashed red line is the baseline fan power estimated by linear interpolation between short windows of data immediately preceding the control period (just prior to 11:00) and just before the next control period (just before 13:00). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

This excess cooling power is absorbed by cooling the building thermal mass. The control is not perfect, and the peak change in fan power consumption approaches ~ 35 kW. At 11:15, the control algorithm is instructed to decrease fan power consumption by 60 kW for another 15 min, i.e. 30 kW below the baseline at 11:00 am. The AC cooling power is lowered below the internal heat generation and heat leak into the building. During this time, the heat leak and internal heat generation are absorbed by the building thermal mass. At 11:30, the open-loop DR control is released, i.e. all thermostat set-points are returned to their nominal values prior to 11:00.

If the processes described above were reversible, the amount of heat extracted from and injected into the building thermal mass during the two control periods would be the same and the building state at 11:30 would be identical to the building state at 11:00. However, the fan power remains elevated for approximately 60 min following the release of the DR control. In effect, the apparently zero-net electrical DR control during the first two control periods between 11:00 and 11:30 has driven the building away from its quasi-steady equilibrium. Additional cooling power (and therefore electrical energy) is required to restore this equilibrium after 11:30.

The behavior of the building and AC-based DR in Fig. 4 is qualitatively similar to the lossy battery charge/discharge analogy discussed in Section 4 and shown in Fig. 2. The thermal energy stored in the building mass is the internal state variable that plays the analogous role to the battery SOC, and variations in the AC fan power around the baseline play the role of battery charge/discharge power. During the first period of increased electricity consumption, the AC system is asked to remove heat from the building at a rate higher than the baseline. During the second control period of decreased electricity consumption, the AC system is asked to remove heat from the building at a rate lower than the baseline. Electrical, hydrodynamic and thermodynamic losses associated with disturbing the system from its quasi steady state subsequently force additional electricity consumption to restore the building's "SOC". Using this analogy, the round-trip efficiency of each of the

DR-AS experiments is computed from the measured data in a similar way as for the battery, i.e.

$$\eta_{RT} = \frac{E_{out}}{E_{in}}, \quad (2)$$

where E_{out} = the total energy below the baseline; E_{in} = the total energy above the baseline, with the energy totals computed over the entire two-hour span of each measurement (see Figs. 3 and 4)

6. Results

6.1. Round trip efficiency – current study

Our experiments consist of repetitions of the DR control pulses shown in Fig. 3 over many days to sample the building performance over a range of conditions. The AC fan power depends on many variables including the outside air temperature, which changes throughout the day. As a result, many of the baselines are not flat but instead slope upwards in the morning hours and downwards in the evenings. For each DR pulse, the round-trip efficiency η_{RT} is computed using (2). The distribution of all the η_{RT} samples is plotted in Fig. 5 (filled circles). The samples of η_{RT} display significant scatter, with a few rare samples giving values greater than one. The variability in η_{RT} is likely because of the variability in the building and AC system conditions that are beyond our control. Averaging over all of the samples, we find $\langle \eta_{RT} \rangle = 0.46$ which is a rather low round-trip efficiency compared to most battery storage systems [11].

As noted in [14], a VAV must provide a minimum ventilation flow to the conditioned space it serves and is restricted in the amount that it is allowed to close. This saturation will affect the AC fan power and may potentially bias measurements of the round-trip efficiency. To test whether our experiments are subject to such bias, DR control pulses of both polarities are used so that the potential VAV saturation is experienced in the first or second control period. The data samples are partitioned into two sets according to the polarity of the pulses. The DR pulses that initially open the

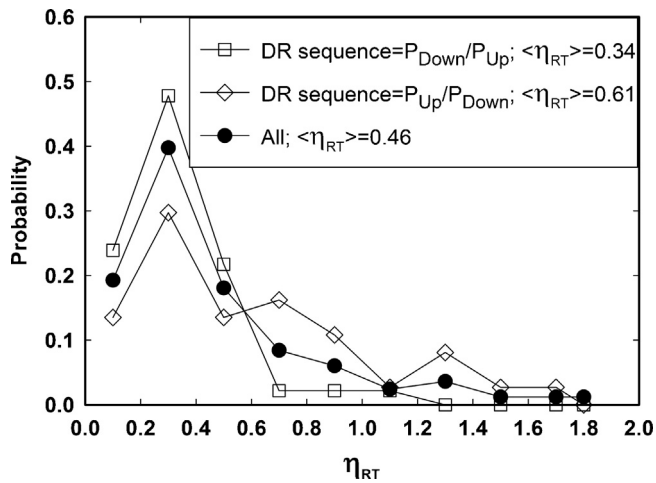


Fig. 5. Probability of observing a round-trip efficiency of η_{RT} for a 30 kW-amplitude DR-AS pulse as shown in Fig. 4. The filled circles show the results for all 83 measurement samples. The open squares show the results for the measurement samples when the AC power was first decreased and then increased (46 samples). The open diamonds display the results for samples when the AC power was first increased and then decreased (37 samples). The measurements are binned in this way because, under the majority of conditions tested in this work, an initial 30-kW decrement in fan power is expected to cause many of the VAV dampers to saturate at their minimum opening causing an asymmetry in the fan power response. An initial increase in fan power is not expected to encounter this initial saturation [14].

damper valves, i.e. increase the fan power, show $\langle \eta_{RT} \rangle = 0.61$ and a larger spread in the measured values as compared to the DR pulses that initially close the damper valves ($\langle \eta_{RT} \rangle = 0.34$). However, this difference does not materially change the discussion in the remainder of this manuscript.

6.2. Round trip efficiency – previous studies

An average round-trip efficiency of 0.46 may appear to be rather low for the AC system used in this work, however, additional analysis of other experiments and simulations reveal similar round-trip efficiencies for space conditioning applications. For example, experiments in [8] utilized a transactive control scheme to control the demand of several types of home appliances including electric space heating and cooling. In contrast to the more complex AC system studied in this work, the residential heating and cooling systems in [8] were far simpler with on/off hysteresis control. Fig. 6 is an estimate of the data from Fig. 7.10a) of [8] that displays the measured average space-conditioning electrical demand versus the time of day and the counterfactual estimate, i.e. an estimate of what the demand would have been if there had been no load control. Lower prices during the nighttime encouraged higher electrical demand for precooling/preheating of the space followed by lower electrical demand during daytime hours when prices were generally higher. Taking the counterfactual as the baseline in Fig. 6, we compute the average energy “delivered to the grid” between the hours 9 and 21 to be 1.1 kW-h and the energy “taken from the grid” between the hours of 0 and 9 and 21 and 24 to be 2.6 kW-h. Assuming that the conditioned spaces end up on average in the same state in hour 24 as they began at hour 0, the round-trip efficiency is 0.42 – a result that is surprisingly close to the $\eta_{RT} = 0.46$ measured in this work. Similar observations have been made in simulations of residential AC systems operated so as to reduce on-peak energy usage [9]. Fig. 10 of [9] shows that moving ~ 4.2 kW-h of energy consumption away from a four-hour-long peak period caused an additional energy consumption of ~ 3.9 kW-h above the baseline. Using these results in (2), we obtain $\eta_{RT} = 4.2 / (4.2 + 3.9) = 0.52$. Although the round-trip

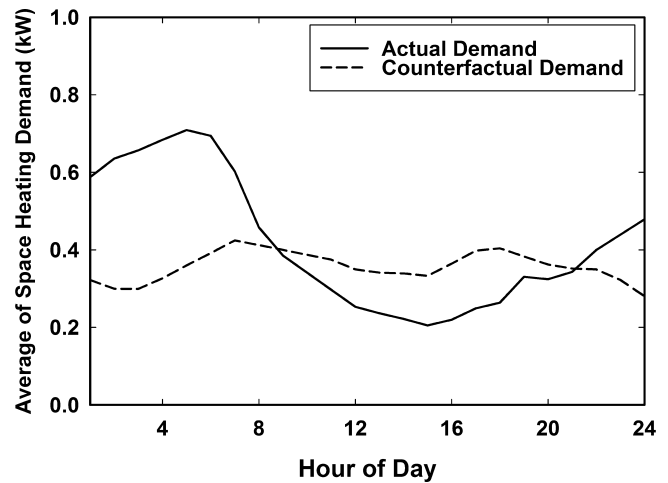


Fig. 6. Adapted from Fig. 7.10a) of [8]. Average space conditioning electrical demand for a range of residences participating in transactive control in a double auction market with five-minute clearing. Both the measured demand (solid line) and counterfactual demand (dashed line) are given. The deviations of the actual around the counterfactual is interpreted as a grid service similar to experimental data in Fig. 5. The round-trip efficiency of this service is estimated to be $\eta_{RT} = 0.42$.

efficiencies in these two cases are similar to that obtained from the measurements in this manuscript, the time scale of the control is significantly longer than for our experiments.

6.3. Effects of chiller – current study

Fan power is only one portion of HVAC electricity consumption. It is important to consider the impact of other power consumptions, e.g. the chillers, on our measurements of round trip efficiency. The experiments in Section 6.1 were reanalyzed to include both fan and chiller power consumption. The example presented in Fig. 7 demonstrates that, similar to the fans, the chiller experiences a post-control period of heightened electricity consumption. Using the same baseline estimation techniques and definition of efficiency as in Section 6.1, we find an aggregate average round trip efficiency of $\langle \eta_{RT} \rangle = 0.42$ indicating that chillers also incur a similar energy penalty while providing DR control.

7. Discussion – economic analysis

The measurements presented above suggest that the participation of large commercial AC systems in fast DR ancillary services may increase total AC energy consumption. In this Section, we discuss the impact of this observation on the economics of fast DR ancillary services from large commercial AC systems, with an emphasis on frequency regulation services.

In many deregulated electrical grids, AS such as frequency regulation are scheduled and paid for via a wholesale ancillary services market. We base the remainder of this discussion on the PJM market [11]. An excellent summary of the PJM frequency regulation market and its relationship to DR from commercial AC systems is given in [4]. Many different types of resources are allowed to bid into the PJM frequency regulation market including generation, grid-level storage, behind-the-meter storage, and demand response. These frequency regulation resources are divided into two groups, i.e. ramp-limited and energy-limited resources. Energy-limited resources can ramp very quickly but cannot sustain an extended request for one sign of regulation. Batteries and flywheels are typical energy-limited resources. Ramp-limited resources can sustain an extended request for one sign of regulation, but are typically mechanically or thermally limited in how fast they can respond. Our

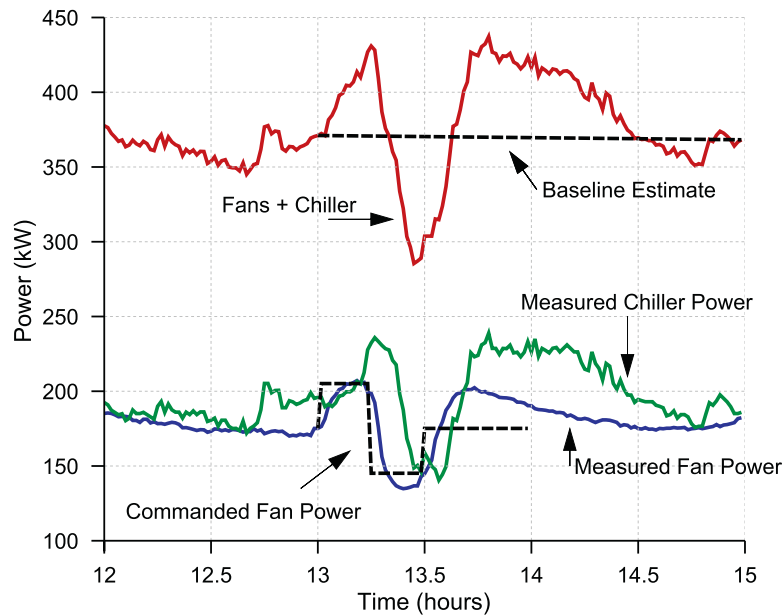


Fig. 7. An experimental fan power trace with chiller power included as well. A square wave pulse is broadcast to the VAVs at 13:00, resulting in the measured fan power deviation shown. The fan power deviation is mirrored in the chiller power consumption, with a time delay induced by the dynamics between the warm return air and cooled water loop. The combined fan + chiller power consumption (the upper red trace) is used to determine efficiency of the overall system. Note the elevated post-control power consumption for the combined power, which is similar to the phenomenon observed when only considering fan power consumption (as in Fig. 4). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

AC-based DR-AS suffers from communication latency within the building combined with a relatively slow VAV mechanical response. Consequently, several minutes of delay can occur in reaching full response (see Fig. 4). Compared to batteries or flywheels, this response is quite slow. Unlike a generator though, our AC-based DR-AS cannot sustain a very long request for one sign of power change. However, on the relatively short timescales of frequency regulation, this is not a serious limitation. We conclude that our AC-based DR-AS fits best into the ramp-limited resource category.

If cleared in the market, a frequency regulation resource responds to the grid operator's automatic generation control (AGC) signal. In PJM, the AGC signal is filtered before it is sent to better accommodate the different types of resources. Energy-limited resources receive a high-pass filtered AGC (RegD signal) that crosses zero very frequently to limit their charge/discharge time. Ramp-limited resources receive a low-pass filtered AGC (RegA signal) that displays a typical ramping time-scale of 5 min. PJM uses test signals [31] that demonstrate the typically expected responses to qualify resources for participation in the frequency regulation market. As currently implemented, the communication latency and relatively slow mechanical response of the VAVs will limit our control to speeds acceptable for RegA service. We will therefore use the PJM RegA market as a basis for the remainder of this discussion.

The frequency regulation market is a bid-based auction-style market in which the marginal unit sets the clearing price for all participants. To participate and compete in this market, resources should submit bids that reflect their cost of providing frequency regulation service. The PJM market [4,11] requires a Cost-Based Regulation Offer that is split into a capability component and a performance component [11] whose definitions are tailored for fuel-consuming generators. Here, we provide our interpretation for a DR-AS regulation resource:

1 Capability Component/Steady Operation Fuel Costs – For a fuel-fired resource, this is the steady-state increase in fuel costs associated with operating the resource at power output that is different from its optimal economic dispatch point. Our DR-AS

resource is not making a steady-state change in its operation and this cost is effectively zero.

- 2 Capability Component/Margin Adder – Up to \$12/MW of margin can be added to the Capability Component. The margin is included at the discretion of the resource owner and will not be considered in this analysis.
- 3 Performance Component/Variable Operation Fuel Costs – For a fuel-fired resource, this is the increase in fuel costs because of non-steady operation. This effect is directly tied to our results for round-trip efficiency in Section 6 and the associated increase in costs because of higher electricity consumption.
- 4 Performance Component/Variable Operation and Maintenance Costs – PJM defines a cap on these costs for different generating unit types. No cap is given for demand resources, and we assume this cap is zero.

From the description of the PJM frequency regulation market above, we conclude that the major component of our resource's DR-AS bid into PJM's RegA market would be the increased electricity costs arising from variable AC operations. However, we must be careful in comparing this cost with historical market data. The model of DR-AS inefficiency described in Section 5 is based on the energy exchanges with the electrical grid above or below baseline operations, while the PJM market assumes that the primary cost driver is an increase in fuel cost for making *changes* in generator output, i.e. regulation mileage [11]. PJM publishes average regulation mileage data, but it does not provide data on the regulation energy exchanges above and below the generator set-points.

Historical data for the PJM RegA signal from December 18, 2012 to January 18, 2014 show that, on average, a 1 MW capacity RegA resource “injects” 0.14 MW-h of energy into the grid and “absorbs” 0.35 MW-h of energy from the grid each hour. Here, the energy “injected” is computed using only times when the RegA signal is positive and “absorbed” is computed using only times when the RegA signal is negative. The bias toward a negative RegA signal is persistent throughout the data investigated and may reflect a bias in the PJM market or reliability operations.

To simplify the market analysis of our DR-AS resource, we shift the zero of the RegA signal to balance the energy exchanges so that, on average, energy $E_{ex} \sim 0.25$ MW-h is injected and absorbed during one hour of frequency regulation service. Under these conditions, the extra energy consumption E_{lost} because of round-trip losses can be computed from (2) as

$$E_{lost} = \frac{1 - \langle \eta_{RT} \rangle}{\langle \eta_{RT} \rangle} E_{ex}. \quad (3)$$

The additional energy required for a DR-AS with $\langle \eta_{RT} \rangle = 0.46$ is 0.29 MW-h for each MW and each hour of frequency regulation service. The cost of this additional energy can vary widely depending on the tariff structure of the utility provider. For an energy tariff of \$0.10/kW-h the cost of providing this DR-AS service is \$29/MW/h which is typically near the marginal cost for this service in the PJM market [32]. However, this cost does not include any profit for the owner (“margin adder”) of the asset providing the DR-AS service.

8. Conclusions and future work

We have developed a methodology for measuring the round-trip efficiency of large commercial air conditioning loads controlled for fast demand response (DR). We have applied this methodology to a large commercial air conditioning system that is controlled on the 15-min time scale – a speed that is approaching that needed for participating in many frequency regulation markets. For the large commercial air conditioning load studied in this work, we find a round-trip efficiency of approximately 0.46 when only considering fan power, and an efficiency of 0.42 when analyzing the combined power of the chiller and fans. These rather low results are very similar to the round-trip efficiencies 0.42 and 0.52 we have estimated from data in other experimental [8] and simulation work [9], respectively. Although the timescales in this other work are quite different, the similarity of the results suggests that the loss mechanisms may be related. It is therefore possible that this low round-trip efficiency may carry over to other space conditioning loads.

There are many important future directions and extensions of this work:

- 1 repeating these measurements on additional building and air conditioning types to explore if the round-trip efficiencies are similar
- 2 repeating these measurements using historical PJM RegA signals to determine the average round-trip efficiency under realistic operation conditions
- 3 performing purposeful experiments on large aggregations of residential space conditioning loads on fast time scales to characterize the losses in these systems
- 4 repeating these experiments on a range of timescales to help identify the source of the losses
- 5 investigation of the causes of the round-trip losses and potential remedies to improve the round-trip efficiency
- 6 characterizing the round-trip efficiency of loads other than space conditioning loads.

Acknowledgements

The authors thank Elaine Hale, Paul Denholm, Marissa Hummon, David Chassin, and Johanna Mathieu for useful discussions. The work at LANL was carried out under the auspices of the National Nuclear Security Administration of the U.S. Department of Energy at Los Alamos National Laboratory under Contract No. DE-AC52-06NA25396. The authors acknowledge the support of the Microgrid Program in the US Department of Energy Office of Electricity.

References

- [1] C. Holly, Renewables to Crush Conventional Generators due to RPS Rules – Study, 2013 October <http://www.theenergydaily.com/publications/ed/Renewables-to-crush-conventional-generators-due-to-RPS-rulesstudy.10217.html>
- [2] C.-J. Yang, R.B. Jackson, Opportunities and barriers to pumped-hydro energy storage in the United States, *Renew. Sustain. Energy Rev.* 15 (1) (2011) 839–844.
- [3] K. Subbarao, J. Fuller, K. Kalsi, R. Pratt, S. Widergren, D. Chassin, Transactive control and coordination of distributed assets for ancillary services. Technical Report PNNL-22942, Pacific Northwest National Laboratory, 2013.
- [4] Peng Zhao, Gregor P. Henze, Sandro Plamp, Vincent J. Cushing, Evaluation of commercial building HVAC systems as frequency regulation providers, *Energy Build.* 67 (2013) 225–235.
- [5] Z. Zhou, F. Zhao, J. Wang, Agent-based electricity market simulation with demand response from commercial buildings, *IEEE Trans. Smart Grid* 2 (4) (2011) 580–588.
- [6] P. Faria, Z. Vale, Demand response in electrical energy supply: an optimal real time pricing approach, *Energy* 36 (8) (2011) 5374–5384.
- [7] N. Lu, An evaluation of the HVAC load potential for providing load balancing service, *IEEE Trans. Smart Grid* 3 (3) (2012) 1263–1270.
- [8] D.J. Hammerstrom, et al., Pacific Northwest GridWise Testbed Demonstration Projects. Part I. Olympic Peninsula Project. Technical Report PNNL-17167, Pacific Northwest National Laboratory, 2007.
- [9] W.J. Cole, K.M. Powell, E.T. Hale, T.F. Edgar, Reduced-order residential home modeling for model predictive control, *Energy Build.* 74 (0) (2014) 69–77.
- [10] D.S. Callaway, Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy, *Energy Convers. Manage.* 50 (5) (2009) 1389–1400.
- [11] PJM, Ancillary Services, 2013, June <http://www.pjm.com/media/training/core-curriculum/ip-gen-301/gen-301-ancillary-services.ashx>
- [12] M. Bourhan Tashtoush, M.M. Al-Rousan, Dynamic model of an HVAC system for control analysis, *Energy* 30 (10) (2005) 1729–1745.
- [13] M.-Y. Chow, J. Teeter, Reduced-order functional link neural network for HVAC thermal system identification and modeling, in: International Conference on Neural Networks, 1997, vol. 1, IEEE, 1997, pp. 5–9.
- [14] G. Goddard, J. Klose, S. Backhaus, Model development and identification for fast demand response in commercial HVAC systems, *IEEE Trans. Smart Grid* 5 (4) (2014) 2084–2092.
- [15] M. Ann Piette, O. Sezgen, D.S. Watson, N. Motegi, C. Shockman, L. Ten Hope, Development and Evaluation of Fully Automated Demand Response in Large Facilities. Technical Report CEC-500-2005-013, Lawrence Berkeley National Laboratory, 2005.
- [16] N. Motegi, M. Ann Piette, D.S. Watson, S. Kiliccote, P. Xu, Introduction to Commercial Building Control Strategies and Techniques for Demand Response. Technical Report LBNL-59975, Lawrence Berkeley National Laboratory, 2007.
- [17] J.H. Eto, J. Nelson-Hoffman, C. Torres, S. Hirth, B. Yinger, J. Kueck, B. Kirby, C. Bernier, R. Wright, A. Barat, et al., Demand Response Spinning Reserve Demonstration. Technical Report LBNL-62761, Lawrence Berkeley National Laboratory, 2007.
- [18] G. Ghatikar, R. Bienert, Smart grid standards and systems interoperability: a precedent with OpenADR, in: Proceedings of the Grid Interop Forum, 2011.
- [19] J.J. Kim, R. Yin, S. Kiliccote, Automated Demand Response Technologies and Demonstration in New York City using OpenADR. Technical Report LBNL-6470E, Lawrence Berkeley National Laboratory, 2013.
- [20] L. Lu, W. Cai, Y. Chai Soh, L. Xie, S. Li, HVAC system optimization—condenser water loop, *Energy Convers. Manage.* 45 (4) (2004) 613–630.
- [21] A. Yuvraj, Building Aggressively Duty-Cycled Platforms to Achieve Energy Efficiency, PhD thesis, University of California at San Diego, La Jolla, CA, USA, 2009, AAI3369624.
- [22] A. Aswani, N. Master, J. Taneja, D. Culler, C. Tomlin, Reducing transient and steady state electricity consumption in HVAC using learning-based model-predictive control, *Proc. the IEEE* 100 (1) (2012) 240–253.
- [23] J. Široký, F. Oldewurtel, J. Cigler, S. Privara, Experimental analysis of model predictive control for an energy efficient building heating system, *Appl. Energy* 88 (9) (2011) 3079–3087.
- [24] J. Ma, J. Qin, T. Salisbury, P. Xu, Demand reduction in building energy systems based on economic model predictive control, *Chem. Eng. Sci.* 67 (1) (2012) 92–100.
- [25] F. Oldewurtel, A. Ulbig, A. Parisio, G. Andersson, M. Morari, Reducing peak electricity demand in building climate control using real-time pricing and model predictive control, in: 49th IEEE Conference on Decision and Control (CDC), IEEE, 2010, pp. 1927–1932.
- [26] J.C. Fuller, N. Prakash Kumar, C.A. Bonebrake, Evaluation of Representative Smart Grid Investment Grant Project Technologies: Demand Response. Technical Report PNNL-20772, Pacific Northwest National Laboratory, 2012.

- [27] H. Hao, A. Kowli, Y. Lin, P. Barooah, S. Meyn, Ancillary service for the grid via control of commercial building HVAC systems, in: American Control Conference (ACC), IEEE, 2013, pp. 467–472.
- [28] J. MacDonald, S. Kiliccote, Commercial Building Loads Providing Ancillary Services in PJM, in: ACEEE Summer Study on Energy Efficiency in Buildings, 2014.
- [29] J. Murphy, High-performance VAV systems, ASHRAE J. (2011, October) 18–28.
- [30] J.L. Mathieu, D.S. Callaway, S. Kiliccote, Variability in automated responses of commercial buildings and industrial facilities to dynamic electricity prices, *Energy Build.* 43 (12) (2011) 3322–3330.
- [31] PJM, RegA Test Wave, 2012, August <http://www.pjm.com//media/markets-ops/ancillary/rega-test-wave.ashx>
- [32] PJM, PJM Regulation Zone Preliminary Billing Data (2013, April) <http://pjm.com/markets-and-operations/market-settlements/preliminary-billing-reports/pjm-reg-data.aspx>