Ancillary services through demand scheduling and control of commercial buildings

Yashen Lin, *Member, IEEE,* Prabir Barooah, *Member, IEEE,* and Johanna L. Mathieu, *Member, IEEE*

***Abstract*—Prior work showed building Heating, Ventilation, Air Conditioning (HVAC) systems can provide ancillary services to the power grid without sacrificing occupant comfort if the reference power variation is of high frequency (seconds to a few minutes). This paper addresses the question of how to do that when the reference power variation is of lower frequency, e.g., periods of a few minutes to an hour. The proposed control system to do so uses a two-layer architecture. An optimizer schedules the baseline cooling and heating power of a building based on load forecasts. A lower level controller is then used to track the scheduled baseline plus ancillary service reference signal. The schedule is periodically updated based on indoor measurements to ensure quality of service in spite of load forecasting error. The algorithm is tested in simulation. Results show that ancillary services in the frequency range of $f \in [1/(1\ \textbf{hour}), 1/(10\ \textbf{minutes})]$ can be extracted from commercial building HVAC systems while still maintaining a comfortable indoor climate.**

***Index Terms*—Ancillary services, demand scheduling, commercial building, HVAC**

Nomenclature

List of Variables

α	Outside air ratio.
δP	Power variation.
ω	Humidity ratio.
h	Specific enthalpy.
m	Flow rate.
P	Power.
S	Score.
T	Temperature.

List of Parameters

Δt_I	Implementation period.
Δt_S	Scheduling period.
γ	Design parameter for schedule update.
λ	Design parameter for schedule update.
τ	Time constant.
C	Heat capacity.
$h_{w,e}$	Latent heat of vaporization of water.
t_d	Transport delay.

List of Superscripts

b	Baseline.
I	Implementation period.
r	Reference.
S	Scheduling period.
s	System operator.

List of Subscripts

a	Air flow.
c	Chiller.
ca	Conditioned air.
cc	Cooling coil.
chw	Chilled water.
d	Disturbance.
f	Supply air fan.
l	Load.
ma	Mixed air.
oa	Outside air.
r	Return.
ra	Return air.
rh	Reheat.
s	Supply.
w	Water.

YL and JM are with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor; yashenl,jlmath@umich.edu. PB is with the Department of Mechanical and Aerospace Engineering, University of Florida, Gainesville; pbarooah@ufl.edu. The majority of this research was conducted when YL was at the University of Florida. This research was supported in part by the NSF grant ECCS-0925534 and the University of Michigan Energy Institute.

I. Introduction

Ancillary services are needed to correct the mismatch between demand and supply in a power grid. Traditionally, ancillary services are provided by generators. An alternative is to use flexible loads which may have less environmental impact and cost in the long run. A number of works have explored the use of residential and commercial building loads to provide ancillary services [1–11].

Ancillary services are distinguished by their time-scales. *Frequency regulation* has a time-scale of a few seconds to a few minutes, while *load following* has a time-scale of minutes to hours [12]. The topic of this paper is the use of loads to provide ancillary services in the time scale similar to load following. Specifically, we propose a method for Variable Air Volume (VAV) Heating, Ventilation, Air Conditioning (HVAC) systems to provide ancillary service in the time scale of ten minutes to an hour. In previous work, it was shown how fan motors in VAV HVAC systems can be used to effectively provide frequency regulation without any discernible variation of indoor climate [7, 8, 13]. In this paper, we extend the time scale and use both chiller and fan motors.

The control system proposed in this paper has two tasks. One is to vary the electric power consumption of the HVAC system in such a way that the deviation from the baseline follows an exogenous reference signal that is band-limited to the above-mentioned timescale: ten minutes to an hour. The baseline is the power the HVAC system would have consumed if it were not providing ancillary service. The other task is to maintain the consumers' Quality of Service (QoS). Specifically, the resulting indoor climate variation should be small enough to be undetectable by occupants, and total energy should not increase significantly.

We assume the system operator broadcasts a signal that contains information on grid-level demand-supply imbalance to all buildings. At each building, this signal is passed through a bandpass filter to generate the reference signal for that building. The bandpass frequencies and bandpass gain of the filter have to be chosen so that the reference signal is limited to a frequency and amplitude that is appropriate for the building's characteristics.

If the reference signal is of high frequency, the baseline can be estimated with a low pass filter since the power variation caused by the climate control system is much slower than that caused by reference tracking. This approach was used in [7, 8, 13]. However, if the reference signal is of low frequency, estimating the baseline is challenging because of lack of time scale separation: the time scales of the power variation caused by the existing climate control system overlaps with that caused by reference tracking. This has been a recognized issue in demand response; see [14] and references therein. We therefore choose to schedule the baseline rather than to estimate it. This approach is inspired by the early work of Borenstein et al. [15] who called it Build Your Own Baseline (BYOB). We therefore call the proposed control system, which includes the baseline scheduler at a supervisory level and an inner loop reference tracking control loop, the BYOB Ancillary Service Controller (BYOB-ASC). In addition, when the reference signal is of high frequency as in [8], variation in the actuation (air flow rate) only affects the fan power but not the chiller power due to the high thermal inertia of the chiller. In the time scale considered in this paper, airflow variations will affect both fan and chiller power, making the design of the control system more challenging.

The control system we present in the paper has the following features. One, it is simple to design and implement. The design requires only a few parameters, black-box Linear Time Invariant (LTI) models of a few HVAC system components, and a few weeks of historical data of HVAC power consumption to fit a load prediction model. The parameters and black-box LTI models can be determined from system identification experiments. The scheduling requires solving a low-dimensional Linear Program (LP) every hour. Two, the system is highly robust to error in load forecasts from weather predictions. This is achieved by the introduction of feedback: the load forecast is updated every hour based on measurements of indoor temperature. As a result, the room temperature deviation is kept small in spite of the forecast error. Theoretical analysis of this robustness is provided.

The three papers most relevant to ours are [9–11], which propose methods to use commercial buildings for frequency regulation (as opposed to load following considered here) including baseline scheduling. This paper makes several contributions compared to [9–11]. The first is the simplicity of our proposed scheduling algorithm. The scheduling in [9–11] requires solving Model Predictive Control (MPC) problems with simplified building thermal dynamical model, while our formulation solves a low-dimensional LP based on steady state load balance. Also, only a thermal capacitance parameter is required within our scheduling problem; thermal resistances are not needed. The second is that [9–11] do not consider the challenges associated with tracking a reference power signal computed by a baseline scheduler, such as transport delay in the water loop. In [9, 10], it is assumed that a heat pump controller can be designed to provide perfect tracking. The control system proposed in this paper includes a low level reference tracking controller that is robust to chilled water loop delay uncertainty. The third contribution of our work is robustness to load forecast error and model mismatch. We explicitly introduce feedback in the scheduling algorithm to correct the forecast error through baseline updates, provide theoretical analysis of the robustness, and evaluate the overall system in simulations that include significant plant-model mismatch. In [9, 10], it is assumed that perfect forecast of building load is available. Ref. [11] includes a disturbance in the building model, but no analysis of its impact is provided. MPC simulations in [9, 10] do not include plant-model mismatch.

On the other hand, [10] considers using a collection of commercial buildings while we consider only one. While [9–11] compute the time-varying reserve capacity, we assume the reserve capacity is constant. Our work is therefore complementary to that in [9–11].

This paper is an extension of the work reported in [16, 17]. In [16], it is assumed that the baseline is somehow known, while both [16, 17] neglected fan power assuming it is insiginifant compared to chiller power, which is not accurate. At the ancillary services timescales we consider, variation in air flow rate will lead to variation in both fan and chiller power[1] and this paper considers both fan and chiller power consumption in the design. Doing so introduces an additional level of complexity as these two sources of demand have distinct dynamics: the relation between air flow rate and fan power is of high frequency with negligible transport delay, while that between air flow and chiller power is of lower frequency with a non-negligible transport delay. The second contribution over the preliminary work [16, 17] is an extensive simulation evaluation of the revised design. We thoroughly examine not only the ancillary service performance but also the impact on QoS (indoor climate and total energy use).

The rest of the paper is organized as follows. Section II describes the proposed system; Section III presents analysis of indoor climate variation and ancillary service capacity; Section IV discusses the simulation setup and results; Section V concludes the paper and discusses directions for future work.

II. The Proposed BYOB-ASC System

We start with a description of conventional chilled-water VAV HVAC systems. The BYOB-ASC leaves certain parts of the climate control system untouched while overriding others.

A. Background - VAV HVAC systems

Figure 1 shows a schematic of a typical single-duct, chilled water-based, VAV HVAC system. Although the method proposed in this paper is applicable to multi-zone HVAC systems with minor modifications, we focus on single zone systems,

[1]The only remaining source of electric power in commercial HVAC system, pumps, are neglected due to their much lower demand. Heating is typically provided by gas, not electricity, hence not considered.

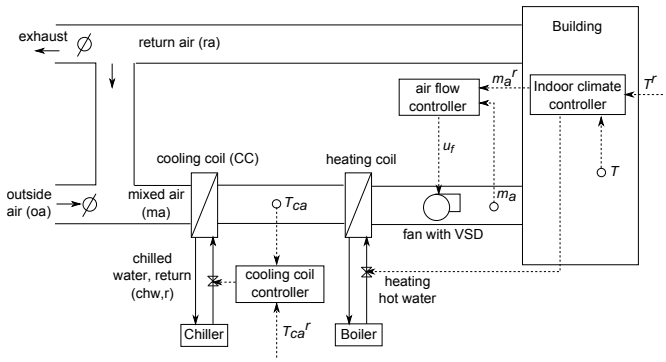

Fig. 1. Schematic of a VAV HVAC system.

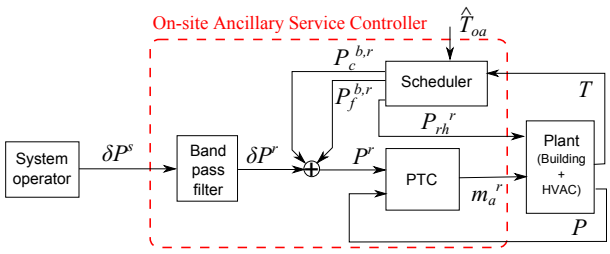

Fig. 2. Schematic illustration of the BYOB-ASC.

such as the one shown in Fig. 1. Fresh outside air (oa) is brought in and mixed with the return air (ra), which ensures good Indoor Air Quality (IAQ). The mixed air (ma) then goes through a cooling coil (CC) in the Air Handling Unit (AHU), where the air is cooled and dehumidified, producing conditioned air (ca). The chilled water (chw) used to condition the air in the cooling coil is produced in a chiller. There may also be a reheat (rh) coil to increase the temperature of the conditioned air. The conditioned air is then distributed to the zone by a supply air fan through ducts. The abbreviations indicated here will be used as subscripts to differentiate different air streams in the following sections.

The indoor climate controller acts on the error between the measured indoor temperature T and the pre-determined set-point T^r to determine the reference m_a^r for airflow rate m_a (kg/s), and reheating rate P_{rh}^r (W). The conditioned air temperature T_{ca} is maintained at a pre-determined set point T_{ca}^r by the cooling coil controller, which varies the chilled water flow rate m_{chw} using valves and chilled water pumps. A controller varies the fan motor speed so that the measured airflow rate m_a tracks the reference m_a^r.

There are three main sources of power consumption in such an HVAC system - mechanical (fan+pumps), cooling (chiller motors), and (re-)heating. In many cases, heating is provided by steam rather than electricity. Moreover, electrical demand of pumps is much smaller than the fans and chillers. Therefore, in this paper we consider *fan and chiller power as the only sources of flexible electrical demand* that can be used for ancillary services. However, when examining the total energy consumption of the HVAC system in Section IV-F, we take into account reheating.

B. The BYOB Ancillary Service Controller

We assume that the system operator broadcasts a signal $\delta P^s(t)$ to all participating resources, such as generators and smart loads. It is argued in [18] that it is imperative to respect the bandwidth limitations of resources in order to maintain consumers' quality of service. Therefore, the proposed system uses an on-site bandpass filter that filters $\delta P^s(t)$ to create a reference power deviation signal, $\delta P^r(t)$, whose magnitude and frequency is appropriate for the building. We also assume that the reference signal the building needs to track is energy neutral.

The advantage of using an on-site filter is that the system operator could utilize the participating buildings to their full potential since the filters would be specifically designed for individual buildings. The disadvantage, however, is that the system operator has to ensure the sum of ancillary service from all buildings meets its total requirements. In this paper, we do not address this problem, and we assume the system operator has enough resources to meet its requirements. Alternatively, the system operator could filter the signal and send different signals to different resources, as the PJM Interconnection (PJM) does [19]. In this case, our proposed approach still works; each building would respond to the signal that best fits their characteristics.

The measured HVAC *electrical* power consumption $P(t)$ is the sum of the fan and chiller power, i.e., the total electrical HVAC power consumption, neglecting the pump power, which is comparatively small. The value of $P(t)$ is the sum of the baseline electrical power $P^b(t)$ and the deviation from it, $P(t) := P^b(t) + \delta P(t)$. In our BYOB approach, the baseline is *scheduled* ahead of time to match the forecasted thermal load. The scheduled baseline power is called $P^{b,r}$. The reference command for the total electrical power consumption is the sum of the scheduled baseline and the ancillary services reference signal:

$$P^r(t) := P^{b,r}(t) + \delta P^r(t) \tag{1}$$

A schematic illustration of the structure of the BYOB-ASC for a building is shown in Fig. 2. The BYOB-ASC consists of two modules: a high-level scheduler and a low-level Power Tracking Controller (PTC). The scheduler computes baseline $P^{b,r}(t)$ ahead of time. The PTC commands the desired airflow rate $m_a^r(t)$ so that the electrical power $P(t)$ tracks the reference $P^r(t)$. We assume that the actual heating power can be made to track the scheduled heating power $P_{rh}^r(t)$ by actuating the heating valve.

Time is divided into a number of slots for scheduling and implementation purposes, with length Δt_S and Δt_I, respectively, where $\Delta t_I < \Delta t_S$. The main steps are:

1) At the beginning of the k-th scheduling period $\mathcal{T}_k^S := (k\Delta t_I,\ k\Delta t_I + \Delta t_S]$, forecast thermal load for this period. The forecasting algorithm is described in Section II-C1.
2) Determine the desired baseline power during $\mathcal{T}_k^S$ which serves the forecast thermal load while ensuring IAQ constraints are satisfied and actuation limits are respected.

3) During the implementation period $\mathcal{T}_k^I := (k\Delta t_I,\ (k+1)\Delta t_I]$, use the PTC to track the reference fan power and chiller power (1) by varying the supply air flow rate.
4) At the end of implementation period, update the load estimate for the next scheduling period $\mathcal{T}_{k+1}^S$, and go back to step 1).

The BYOB-ASC over-rides the climate control system, i.e., the references $m_a^r(t)$ and $P_{rh}^r(t)$ that are otherwise computed by the climate control system are now computed by the BYOB-ASC. All other local control loops in the building, such as control loops in the cooling coil, heating coil, fan speed controller, are left untouched.

C. Baseline scheduling

1) Thermal load forecasting model: We first define the term "thermal load" precisely and describe how to estimate it so that these estimates can be used to fit a model that is used later for prediction.

Assuming that the existing climate control system is able to maintain indoor temperature perfectly at its set point, the following steady-state relationship will hold:

$$0 = -P_{cc}(t) + P_{rh}(t) + P_l(t) \tag{2}$$

where $P_{rh}(t), P_{cc}(t)$ are the rate of heat provided to and extracted from the building, respectively, by the heating and cooling coils, and the last term $P_l(t)$ is the rate of heat entering the building from all other sources, such as from outside air, solar irradiation, occupants and plug loads. We define the term $P_l(t)$ as the *thermal load* of the building. Equation (2) gives us an estimate of $P_l(t)$ from measurements of $P_{cc}(t)$ and $P_{rh}(t)$, which are obtained from meters installed in the cooling and heating coils.

Once the building's thermal load is estimated, we fit a model to these data. We use the linear-regression model developed in [20], where the load is a piecewise linear function of outside air temperature and of time-of-week. Time of week acts as a proxy for occupant behavior and appliance use. The ordinary least squares method is used to parameterize the model with historical load/temperature measurements. The model is then used to forecast the building's load at a particular time-of-week given a forecast of outside air temperature.

2) Baseline scheduling at $t = 0$: The algorithm for updating the baseline at subsequent scheduling intervals builds on the algorithm used at the first scheduling period $\mathcal{T}_0^S = [0,\ \Delta t_S]$, so we first describe that in detail.

The HVAC system is required to bring in certain amount of fresh outdoor air to ensure IAQ. This is achieved by setting a minimum supply air flow rate, which is computed based on floorspace and occupancy according to [21]. At the same time, to ensure the humidity requirement, the supply air is cooled to a low temperature (55 ^{o}F in this paper, which is widely used in practice) first, no matter what the thermal load is in the building. Under this mechanism, there is effectively a minimum cooling load the HVAC has to provide at any time. When the actual thermal load of the building is lower than this minimum cooling load, reheating has to occur to maintain the comfortable room temperature. This leads to simultaneous cooling and heating. We emphasize that simultaneous cooling and heating occurs in practice quite often for the same reason: to ensure ventilation and humidity constraints.

The objective of the scheduling is to minimize the sum of fan, cooling, and heating energy consumption subject to the load balance, actuator constraints, and IAQ constraints. We formulate the optimization problem as: for $t \in \mathcal{T}_0^S = [0, \Delta t_S]$,

$$\underset{m_a(t), P_{rh}(t)}{\text{minimize}} \int_{\mathcal{T}_0^S} (P_{cc}(t) + \tilde{P}_f(t) + P_{rh}(t))dt \tag{3}$$

subject to

$$-P_{cc}(t) + P_{rh}(t) + \hat{P}_l(t) = 0 \tag{4}$$

$$m_a(t) \in [m_{lb}, m_{ub}], \quad P_{rh}(t) \in [P_{lb}, P_{ub}] \tag{5}$$

where m_{lb} and m_{ub} are the bounds for supply air flow rate $m_a(t)$; P_{lb} and P_{ub} are the bounds for reheat power $P_{rh}(t)$; $\hat{P}_l(t)$ is the forecast of load from the load forecasting model described in Section II-C1.

The cooling power $P_{cc}(t)$ is:

$$P_{cc}(t) = m_a(t)(h_{ma}(t) - h_{ca}(t)) \tag{6}$$

where

$$h_{ma}(t) = C_{p,a}T_{ma}(t) + w_{ma}(t)(h_{w,e} + C_{p,w}T_{ma}(t)) \tag{7}$$

$$h_{ca}(t) = C_{p,a}T_{ca}(t) + w_{ca}(t)(h_{w,e} + C_{p,w}T_{ca}(t)) \tag{8}$$

$$T_{ma}(t) = \alpha(t)T_{oa}(t) + (1-\alpha(t))T(t) \tag{9}$$

$$w_{ma}(t) = \alpha(t)w_{oa}(t) + (1-\alpha(t))w(t) \tag{10}$$

where $w_{*a}(t)$ is the humidity ratio of the air, $C_{p,a}$ and $C_{p,w}$ are the specific heat capacity of air and water, $h_{w,e}$ is the latent heat of vaporization of water at 0^oC, $\alpha(t)$ is the ratio of outside air to supply air. We assume that (i) the cooling coil control loop maintains the conditioned air at its temperature set-point T_{ca}^r; (ii) conditioned air humidity is constant; (iii) $\alpha(t)$, the room temperature $T(t)$, and humidity $\omega(t)$ stay at their set-points, and all the set-points are constants.

The fan power $P_f(t)$ is:

$$P_f(t) = c_f m_a(t)^3 \tag{11}$$

where c_f is the fan power coefficient. To keep the optimization problem convex, we use a linear approximation of the fan power $\tilde{P}_f(t)$ in the objective function:

$$\tilde{P}_f(t) = 3c_f m_{a,*}^2 m_a(t) \tag{12}$$

where $m_{a,*}$ is the nominal supply air flow rate.

Once the optimization problem is solved, we obtain the scheduled supply air flow rate $m_a^{b,r}(t)$. The scheduled fan power $P_f^{b,r}(t)$ is computed by (11). The scheduled baseline cooling coil power reference $P_{cc}^{b,r}(t)$ is obtained from (6). The corresponding power reference for the chiller $P_c^{b,r}(t)$ is then calculated from $P_{cc}^{b,r}(t)$ and the chiller model, which will be discussed in Section II-D. The scheduled baseline reference power for the PTC is simply $P^{b,r}(t) = P_c^{b,r}(t) + P_f^{b,r}(t)$.

Under the assumptions listed above, the specific enthalpies do not depend on the actuation $m_a(t)$, so energy minimization is equivalent to power minimization at every instant over the scheduling period. Therefore, the baseline power schedule is

computed by solving the following optimization problem for every $t \in \mathcal{T}_0^S$:

$$m^r(t), P_{rh}^r(t) = \underset{m_a(t), P_{rh}(t)}{\arg\min} \; (P_{cc}(t) + P_c(t) + P_{rh}(t)) \quad (13)$$

subject to the load and actuator constraints (4) and (5). This is a convex problem, in particular, a linear program in two decision variables, $m_a(t)$ and $P_{rh}(t)$, and can be solved easily. There are multiple sets of $m_a(t)$ and $P_{rh}(t)$ that satisfy the constraints, so the optimization problem has a non-trivial feasible set.

Note that, in our approach, we schedule the HVAC system to meet the load estimated from power consumption data under normal building operation. This maintains the temperature around its normal setpoint. Unlike [9–11], we only harness the flexibility under this normal operation to provide ancillary services. We do not attempt to manipulate temperature setpoints to minimize energy consumption or shift loads between hours, thus allowing us to decouple time periods. As the trade-off for this simplification, the energy consumption of our algorithm could be suboptimal compared to that of the algorithms in [9–11], where the temperature setpoint is free to vary.

Under extreme cases where the thermal load is outside the designed capacity of the HVAC system, the problem (13) could be infeasible. In the case of load higher than the designed capacity range, the maximum supply air flow rate and no heating will be scheduled; in the case of load lower than the designed capacity range, the minimum supply air flow rate and maximum heating will be scheduled.

3) Baseline update: The scheduled baseline will be inaccurate because of forecasting error and imperfect tracking performance of local control loops. In order to prevent the indoor temperature from deviating too far from the setpoint, the baseline is updated periodically based on measured indoor temperature.

Specifically, the baseline is updated at any scheduling period $\mathcal{T}_k^S$, $k \geq 1$ by solving the optimization problem (13), with only one difference: the load constraint (4) is replaced by:

$$-P_{cc}(t) + P_{rh}(t) + \hat{P}_l(t) + P_{c1}^k(t) + P_{c2}^k(t) = 0, \quad t \in \mathcal{T}_k^S \quad (14)$$

where $P_{c1}^k(t), P_{c2}^k(t)$ are two correction terms:

$$P_{c1}^k(t) := \gamma_1 \dot{T}(k\Delta t_I) e^{-\lambda(t - k\Delta t_I)}, \quad t > k\Delta t_I \quad (15)$$

$$P_{c2}^k(t) := \begin{cases} \gamma_2 E_{end} & t \in (k\Delta t_I, \; k\Delta t_I + \frac{1}{\gamma_2}] \\ 0 & t \in (k\Delta t_I + \frac{1}{\gamma_2}, k\Delta t_I + \Delta t_S] \end{cases} \quad (16)$$

where γ_1, γ_2 and λ are design parameters, and $E_{end} := C(T(k\Delta t_I) - T^r)$, where C is the thermal capacitance of the building. In effect, the predicted load is updated from the open-loop forecast by P-D (proportional derivative) feedback. The term $\dot{T}(k\Delta t_I)$ can be causally estimated from measurements of T. Note that for the first scheduling period, the correction terms are zero.

Once the schedule for airflow rate and reheat power are determined, the baseline electrical power reference is computed as described in Section II-C2.

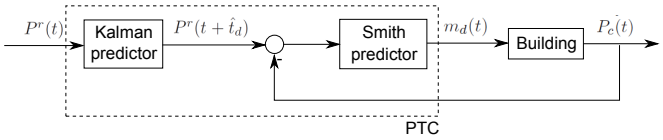

Fig. 3. Architecture of the power tracking controller.

D. PTC design

The objective of the PTC is to make $P(t)$ track the reference signal $P^r(t)$. Power is consumed by the chiller to produce chilled water at temperature $T_{chw,s}(t)$ from the warm chilled water returned from the cooling coil at temperature $T_{chw,r}(t)$. The rate of heat transfered to the chilled water at the cooling coil is $P_{cc}(t) = m_{chw}(t)(T_{chw,r}(t) - T_{chw,s}(t))$, which must be removed by the chiller. Note also that due to conservation of energy, this formula must yield the same result as (6). The transfer function between power demand at the cooling coil and power consumption at the chiller is modeled by a first order lag with a pure time delay:

$$\frac{P_c(s)}{P_{cc}(s)} = \frac{1}{k_c} \frac{1/\tau_c}{s + 1/\tau_c} e^{-t_d s} \quad (17)$$

where t_d is the transport delay arising from the finite speed of water flow from the cooling coil to the chiller, and k_c is the coefficient of performance (COP) of the chiller. This arguably simple model is inspired by the experimental results presented in [22], which shows that step response of a chiller is approximately that of a first order system. Although the COP varies as a function of operating conditions, we assume it is a constant equal to the seasonal average.

The PTC shown in Fig. 3 has two main parts: a Smith predictor to achieve reference tracking with a delay and a Kalman predictor to predict the future reference. The Smith predictor is a classical tool for designing compensators for LTI plants with delay [23]. The Smith predictor ensures that the closed loop transfer function with the delay t_d in the plant is stable when the delay is known, but does not ensure good tracking.

If the reference signal can be predicted t_d time-steps into the future, reference tracking can be achieved by passing the predicted reference to the Smith predictor-based control system. The reference signal $P^r(t)$ consists of two parts; see (1). The scheduled baseline $P^{b,r}(t)$ is known ahead of time, so advancing it by $\hat{t}_d$ is straightforward. The ancillary service reference signal $\delta P^r(t)$, however, is calculated from the signal received from the system operator in real-time and is therefore not known ahead of time. We use a Kalman predictor to predict this part of reference signal [24]. A double integrator model is used for the Kalman filter:

$$\dot{\eta}_1(t) = \eta_2(t), \qquad \dot{\eta}_2(t) = \xi(t), \qquad \delta P^r = \eta_1(t) + v(t)$$

where $\xi(t), v(t)$ are Gaussian white noise. Since the reference signal is smooth, it changes at an approximately constant rate during short time intervals, which motivates the double integrator model. More detail about the Kalman predictor can be found in [25].

Note that in designing the controller, we assume that both fan and chiller power has a transport delay, while in reality only the chiller power has a delay.

III. Temperature dynamics and ancillary services capacity

A. Analysis of the indoor temperature dynamics

The baseline update is designed to create feedback that provides robustness of zone temperature control to load forecasting error and other disturbances. Feedback can, however, cause instability. Here we analyze the effect of that feedback on the temperature deviation dynamics, with load forecasting error and the ancillary service reference signal as exogenous inputs, and the zone temperature deviation from the setpoint as the output. The analysis will provide guidelines on how to choose the design parameters in the baseline update so that the effect of load forecasting error on the temperature deviation is bounded and small. The analysis also provides a means to estimate the capacity of ancillary services the building can provide.

Define $x_1(t) = T(t) - T^r$, and $x_2(t) = \dot{T}(t)$. Note that $\dot{x}_1(t) = x_2(t)$. Consider the k-th implementation period $t \in \mathcal{T}_k^I = (k\Delta t_I,\ (k+1)\Delta t_I]$. With (15) and (16), we have

$$P_{c1}^k(t) = \gamma_1 x_2(k\Delta t_I) e^{-\lambda(t-k\Delta t_I)} \tag{18}$$

$$P_{c2}^k(t) = \gamma_2 C x_1(k\Delta t_I) \tag{19}$$

We assume $\frac{1}{\gamma 2} \geq \Delta t_I$ so that the second correction term $P_{c2}^k(t)$ remains constant during $\mathcal{T}_k^I$, which simplifies the expressions.

Consider the whole building as a single capacitor C, and let $T(t)$ be the indoor temperature. The temperature dynamics of the building can be approximated by:

$$C\dot{T}(t) = -P_{cc}(t) + P_{rh}(t) + P_l(t) \tag{20}$$

Both the load forecast error and the ancillary service reference signal enter the system as disturbances. Let $\delta P_d(t)$ be the disturbance to the system: $\delta P_d(t) := \delta P_l(t) + \delta P^r(t)$, where $\delta P_l(t)$ is the load forecast error: $\delta P_l(t) := P_l(t) - \hat{P}_l(t)$. Equation (20) can be rewritten as

$$C\dot{T}(t) = -P_{cc}^r(t) + P_{rh}^r(t) + \hat{P}_l(t) + \delta P_d(t) \tag{21}$$

Since the scheduled $P_{cc}^r(t)$ and $P_{rh}^r(t)$ have to satisfy (14), we have

$$C\dot{T}(t) = -P_{c1}^k(t) - P_{c2}^k(t) + \delta P_d(t) \tag{22}$$

This leads to

$$x_2(t) = \dot{T}(t) = \frac{1}{C}(-P_{c1}^k(t) - P_{c2}^k(t) + \delta P_d(t)) \tag{23}$$

We now define the states:

$$x_1^k = x_1(k\Delta t_I) \qquad x_2^k = x_2(k\Delta t_I) \tag{24}$$

and denote $\delta P_d^k := \delta P_d(k\Delta_I)$.

It can be shown with straightforward manipulation that the dynamics of the discrete-time states are given by [17]

$$\begin{bmatrix} x_1^{k+1} \\ x_2^{k+1} \end{bmatrix} = \begin{bmatrix} 1-\gamma_2\Delta t_I & \frac{\gamma_1}{C\lambda}(e^{-\lambda\Delta t_I}-1) \\ -\gamma_2 & \frac{\gamma_1}{C}e^{-\lambda\Delta t_I} \end{bmatrix} \begin{bmatrix} x_1^k \\ x_2^k \end{bmatrix} + \frac{1}{C}\begin{bmatrix} \delta W^{k+1} \\ \delta P_d^{k+1} \end{bmatrix} \tag{25}$$

where

$$\delta W^{k+1} := \int_{k\Delta t_I}^{(k+1)\Delta t_I} \delta P_d(t) dt \tag{26}$$

We make a further simplification by assuming that the disturbance is constant during an implementation period, which makes $\delta W^{k+1} = \Delta t_I \delta P_d^k$. The system (25) now also defines a transfer function from the single input δP_d^k (disturbance) to the output x_1^k (zone temperature), which we denote by $G_{TQ}(z)$.

The parameters of the correction terms in the baseline update have to be chosen to make (25) asymptotically stable. In addition, the $\mathcal{H}_\infty$ norm of the transfer function $G_{TQ}(z)$, defined as $\|G_{TQ}\|_\infty := \max_\omega |G_{TQ}(e^{j\omega})|$ should be small, which ensures small deviation of the space temperature from the setpoint.

B. Ancillary service capacity

We define the capacity of ancillary services a building can provide as the amplitude of the ancillary service reference signal δP^r the building can track without significantly affecting the indoor climate. Temperature, humidity, and IAQ are all important attributes of indoor climate. Let $\Delta\bar{T}$ be the maximum allowable variation in room temperature from its setpoint, and let $\Delta\bar{m}$ be the maximum allowable variation in supply air flow rate from its nominal value. An appropriate value of $\Delta\bar{T}$ can be specified by the building operator, and can be obtained from ASHRAE mandated thermal comfort envelope [26]. Since the conditioned air temperature is assumed to be maintained at 55°F by the cooling coil controller, indoor humidity will be maintained as long as sufficient air flow is maintained. The airflow variation bound $\Delta\bar{m}$ comes from ventilation constraints, which in turn ensures adequate IAQ [21].

Let P_{bid} be the capacity of ancillary service the building is willing to provide, which is communicated to the system operator through a market bid or other mechanisms. To keep the indoor temperature within $\Delta\bar{T}$ of its setpoint, we must have

$$(P_{bid} + \delta P_l)\, \|G_{TQ}\|_\infty \leq \Delta\bar{T} \tag{27}$$

The parameter P_{bid} should be chosen to satisfy this relationship. Another constraint on P_{bid} comes from the bound on airflow variation. Let $G_{mp}(s)$ be the closed loop transfer function from the ancillary service reference signal to the variation in supply air flow rate. $G_{mp}(s)$ can be obtained from linearizing a calibrated model of the building and its HVAC system (such as the one used in the simulations in Section IV), or from a system identification experiment. Let $\|G_{mp}\|_f$ be the

largest gain in the frequency range of interest of the reference signal. We must have:

$$P_{bid}\|G_{mp}\|_f \leq \Delta\bar{m} \tag{28}$$

Thus, the maximum reserve capacity $\bar{P}$ the building could provide can be calculated by:

$$\bar{P} = \min\left(\frac{\Delta\bar{T}}{\|G_{TQ}\|_\infty} - \delta P_l,\ \frac{\Delta\bar{m}}{\|G_{mp}\|_f}\right) \tag{29}$$

With the parameters used in the simulation, $\bar{P}$ is calculated to be 5.4 kW.

In [9–11], the temperature setpoint is free to vary within the comfortable range and the reference signal is not necessarily zero-mean. Thus, the reserve capacity the building can provide is time varying. In our approach, we keep the temperature setpoint constant and provide ancillary services only. Thus, the reserve capacity in our approach is constant.

IV. Simulation Study

A. Reference signal

The objective of our control algorithm is to provide ancillary services with frequency ranges suitable for the building HVAC system. For demonstration, we use the Area Control Error (ACE) signal as the signal received from the system operator $\delta P^s(t)$ in the simulation, since it contains information of real-time supply-demand mismatch. The reference signal $\delta P^r(t)$ is obtained by bandpass filtering the ACE with a Butterworth filter whose passband is $f \in [1/(1\text{ hour}), 1/(10\text{ minutes})]$ and passband gain is 5.5×10^{-6}. The resulting maximum magnitude of δP^r is 5 kW, while the total power consumption of the HVAC system is around 10 kW (observed from Fig. 6).

The rationale for choosing this "filtered ACE" as the on-site reference signal is that PJM's RegA and RegD signals, which are broadcasted to ancillary service providers, are computed by filtering the ACE signal [19]. Spectral analysis of RegA and RegD showed that both are higher frequency than the frequency band of interest of this work.

B. Building and HVAC system model

A high-fidelity non-linear dynamic model of a building and its HVAC system is used for simulations. The model is calibrated to data from a large zone in Pugh Hall on the University of Florida campus. The zone is serviced by a dedicated AHU. A brief description of each component of the HVAC system is included below.

A schematic of the components of the plant model ("building+HVAC" in Fig. 2) is shown in Fig. 4. The cooling coil model, closed loop fan control model, and building thermal model (and their calibration) are described in [25], so we omit the details here. The cooling coil controller is a PID controller. The two controllable inputs $m_a^{b,r}(t), P_{rh}^r(t)$ are computed by the BYOB-ASC. Among the exogenous inputs, the weather related inputs ($T_{oa}(t)$, $w_{oa}(t)$) are obtained from weather data, the setpoints of the lower-level control loops (α, $T_{chw,s}$, T_{ca}^r) are set to constant values observed from Pugh Hall's

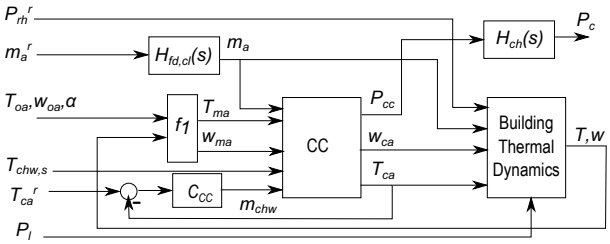

Fig. 4. Schematic of the "plant" (see Fig. 2) used in simulations. CC stands for cooling coil, C_{CC} for cooling coil controller, $H_{fd,cl}$ is the closed loop fan and duct dynamics with the fan controller, H_{ch} is the chiller dynamics. The function f_1 is given by (9) and (10).

Building Automation System (BAS). The thermal load $P_l(t)$ is calculated from Pugh Hall data using (2).

The time constant of the chiller model (17) is taken to be 200 seconds, which is based on experimental results reported in [22]. The transport delay t_d is taken as 30 seconds, which is approximately the time taken by chilled water to travel from a roof-top chiller to the cooling coil in Pugh Hall (computed based on the estimated chilled water flow speed). The COP of the chiller is taken as 3.5 [27].

We emphasize that there is significant model mismatch between the plant and the models used within the control system. For example, in the plant, the building temperature dynamics are modeled with a three-state nonlinear model, while the control system uses a first-order LTI model. Additionally, the scheduler assumes that most air and water related variables remain exactly at their setpoints at all times, while, in the plant, they vary.

C. Ancillary Service Controller

Recall that the baseline scheduler needs the data-driven load model (2). Historical data of $P_{cc}(t)$ and $P_{rh}(t)$ are collected from sensors installed in Pugh Hall. Weather data are collected from *www.wunderground.com* for Gainesville, FL. Two weeks of data (01/13/2014 to 01/26/2014) are used in estimating the parameters of the model. The BYOB-ASC uses this model and 24-hour regional weather forecasts from *www.wunderground.com* to predict the thermal load on the building using (2).

The prediction from the load model is compared with the measured load for the week (04/14/2014 to 04/20/2014). Figure 5 shows the model calibration (in sample) and validation (out of sample) results. The validation results shows that prediction error is non-negligible. The BYOB-ASC is robust to large prediction errors, which the simulation results presented in the next section will show.

The other parameters needed by the baseline scheduler and updater, and the values used in the simulation, are shown in Table I. The compensator within the Smith predictor in the PTC loop is designed so that the sensitivity function of the closed loop system (P^r to P_c) is approximately $\frac{s}{s+0.1}$, which makes the tracking error less than -3 dB when $\omega < 0.1$ rad/s. The models of the components of the HVAC system

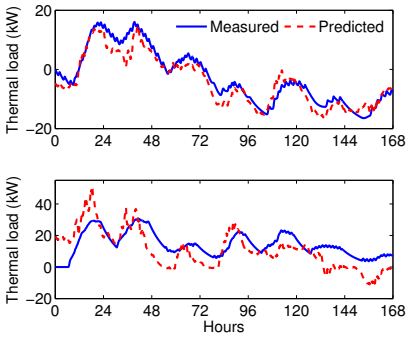

Fig. 5. Thermal load model calibration and validation. Top: measured and predicted thermal load with calibration data 01/13/2014 to 01/19/2014. Bottom: measured and predicted thermal load with validation data 04/14/2014 to 04/20/2014.

TABLE I
PARAMETERS USED BY THE CONTROL SYSTEM.

Parameter	Value	Parameter	Value
C	8×10^7 J/K	γ_1	3.2×10^6 J/K
Δt_I	1 hour	γ_2	1/(1 hour)
Δt_S	24 hours	λ	1/(1 hour)
T^r	72 ^oF	T^r_{ca}	55 ^oF
$T_{chw,s}$	44 ^oF		

described in Section IV-B are connected together to get the full model. The system is linearized to obtain the relevant transfer functions for design of the compensator.

D. Performance metrics

Metrics to quantify performance are described here. Let $\delta P(t) \triangleq P(t) - P^{b,r}(t)$ be the measured power deviation from the scheduled baseline, i.e., the actual ancillary service provided. We quantify the tracking error $e(t) := \delta P^r(t) - \delta P(t)$ by the metric

$$r_R := \frac{\sqrt{\frac{1}{t_0}\int_0^{t_0} e^2(t)}}{\text{avg}_t|\delta P^r(t)|} \tag{30}$$

over an interval t_0.

We also evaluate the tracking performance by the following three scores: S_c - the correlation score, S_d - the delay score, and S_p - the precision score. These criteria are inspired by the PJM scores [19]. Although the ancillary services considered in this paper are different from frequency regulation which the PJM scores are designed for, the criteria evaluate the precision and response speed of the tracking. Define the correlation coefficient to be:

$$R_P(\tau) = \frac{\text{cov}(\delta P(t), \delta P(t+\tau))}{\sigma_{\delta P(t)}\sigma_{\delta P(t+\tau)}} \tag{31}$$

where σ is the standard deviation of the signal. The parameter τ^* is defined as the time shift for which the response has the highest correlation with the reference signal:

$$\tau^* = \underset{\tau \in [0, 5\,\text{mins}]}{\arg\max}\ R_P(\tau) \tag{32}$$

TABLE II
PERFORMANCE VARIATION WITH $\hat{t}_d$ ($t_d = 30$ SEC.)

Test	$\hat{t}_d$	r_R	S_c	S_d	S_p
Filtered ACE (Summer day)	30	0.446	0.959	0.847	0.782
	0	0.609	0.922	0.753	0.567
	20	0.422	0.960	0.860	0.793
	40	0.506	0.939	0.847	0.759
Filtered ACE (Winter day)	30	0.574	0.924	0.793	0.705
	0	0.912	0.828	0.647	0.289
	20	0.574	0.922	0.793	0.709
	40	0.573	0.924	0.793	0.701

The scores S_c and S_d are then determined as:

$$S_c = R_P(\tau^*),\ S_d = \left|\frac{\tau^* - 5\,\text{mins}}{5\,\text{mins}}\right| \tag{33}$$

The precision score S_p is defined as:

$$S_p = 1 - \frac{1}{n}\sum_{i=1}^{n}\frac{|\delta P(i) - \delta P^r(i)|}{|\delta P^{r,a}|} \tag{34}$$

where $\delta P^{r,a}$ is the hourly average of the reference signal and n is the number of samples.

For indoor climate quality, we use the temperature violation D_T defined in [28]. This score is 0 if the indoor temperature stays within the minimum and maximum allowed values, which we set to $70°F$ and $75°F$ [26, Chapter 8]. If the temperature deviates outside of these bounds, the violation score becomes non-zero.

Variation in supply air flow rate affects ventilation, thus smaller variation is preferable. This variation is quantified by $\delta m_{a,avg} := \frac{1}{t_0}\int_0^{t_0}\left|\frac{m_a(t) - m_a^r(t)}{m_a^r(t)}\right| dt$ over an interval t_0.

E. Results

The simulations are conducted in Simulink®. We present simulation results for a summer day and a winter day. During the summer day, cooling load is the dominant load and during the winter day, heating load is the dominant load. The reference signal along with the resulting fan, cooling, and reheating schedules are shown in the top plots in Fig. 6. The tracking performance is shown in the bottom two plots in Fig. 6. The effect of the controller on room climate is shown in Fig. 7. The ancillary services provided by the fan and chiller are shown in Fig. 8. The performance scores are shown in Table II. The results presented in the figures are for the cases when the true value of the chilled water loop transport delay was used in the Smith predictor of the PTC.

On both the summer day and winter day, the room temperature remains in the comfortable range, so the temperature violation D_T is 0. The mean variation $\delta m_{a,avg}$ in supply air flow rate is 25.4% for the summer day simulation and 21.4% for the winter day simulation.

For the values used by the control algorithm in the simulation, the eigenvalues of the state matrix in (25) are 0.52 and 0.04, which satisfies the asymptotic stability requirement of the baseline update-driven temperature dynamics. For these parameter values, $\|G_{TQ}\|_\infty = 4.75 \times 10^{-5}$ (Kelvin/Watt).

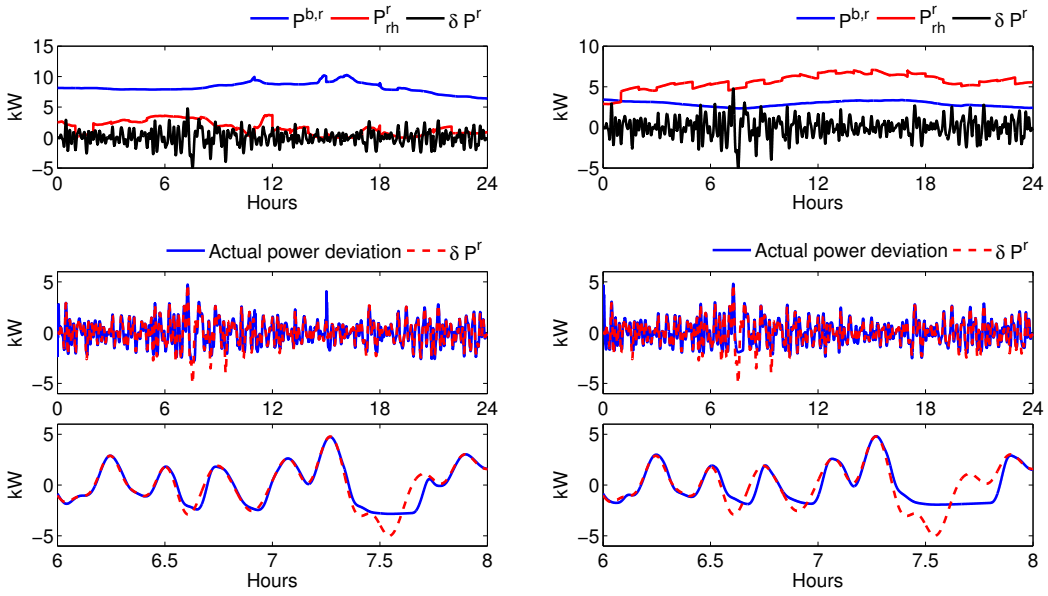

Fig. 6. Simulation results for a summer day (left) and a winter day (right). Top: reference signals $P^{b,r}$, P^r_{rh} and δP^r, which are calculated and updated online. Middle: power deviation tracking performance. Bottom: 2-hour close-up of the middle plot.

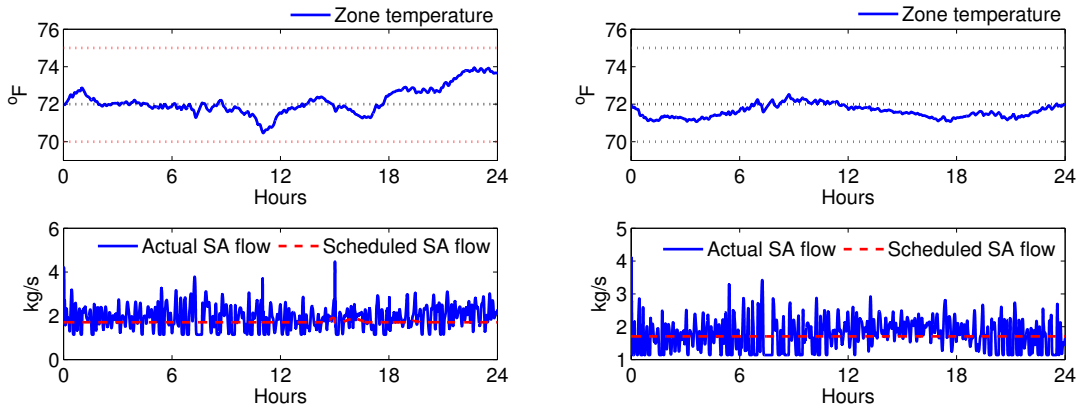

Fig. 7. Room temperature and supply air flow rate during simulation for a summer day (left) and a winter day (right). The red horizontal lines in top figures indicate the comfortable temperature range. The black horizontal line in top figures indicates the temperature setpoint. SA refers to supply air.

The maximum load prediction error is 18 kW (see Fig. 5), and the maximum magnitude of the reference signal is 5 kW. Thus the maximum magnitude of the load deviation from the baseline load is 23 kW, which predicts the maximum deviation of the temperature from the baseline setpoint to be $4.75\times10^{-5}\times23000 \approx 1.1$ Kelvin, or, 2.0 ^{o}F. This prediction is essentially equivalent to the observed value of 2 ^{o}F; cf Fig. 7. One should note that this level of variation around the setpoint occurs during normal operation; see Fig. 8 in [13].

To study how sensitive closed loop performance is to the knowledge of the transport delay in the chilled water loop, we performed simulations with several distinct delay values in the controller that are different from the true delay in the plant. The results are shown in Table II. We see that the control system can handle a ±10 second (33%) error in the estimate of the delay without a significant performance degradation, when the actual delay is 30 seconds. Beyond that, its performance deteriorates significantly.

F. Impact on energy consumption

Energy consumption is another important aspect of QoS of buildings. The authors of [29] reported inefficiencies associated with providing ancillary services from large commercial buildings, which leads to additional energy consumption. In this paper, we consider two types of power consumption changes: (i) variation from the scheduled baseline due to the PTC; (ii) variation from what a conventional climate controller would consume due to the baseline scheduling algorithm.

For (i), we performed simulations with the same scheduled fan, chiller, and reheat power with and without providing ancillary services. The comparison is shown in Fig. 9 and Table III. The difference in total energy consumption during

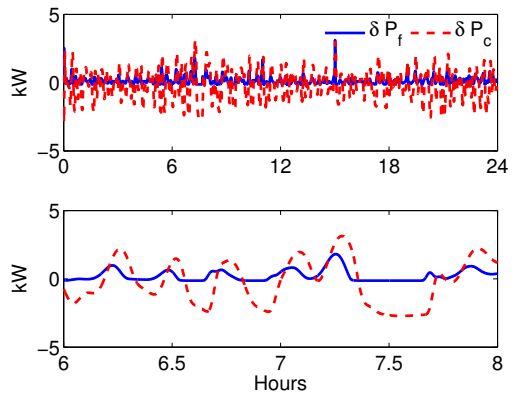

Fig. 8. Comparison between ancillary services provided by the fan and chiller during the summer simulation. Top: power deviations by the fan δP_f and chiller δP_c; bottom: 2-hour close-up of the top plot.

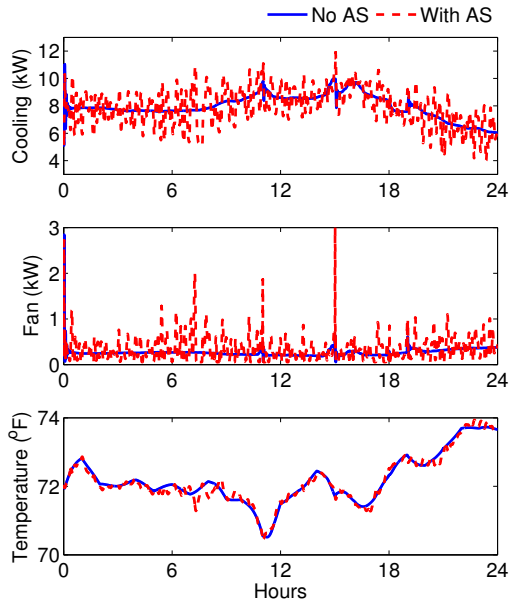

Fig. 9. Comparison between simulations with and without providing ancillary services (AS) during the summer simulation.

the 24-hour simulation period between the two cases is less than 1%. Another observation from these results is that the impact of the ancillary services reference signal on the room temperature is small. Thus, the temperature difference from the setpoint at the end of the simulation observed in Fig. 7 is likely caused by thermal load forecast error, not the ancillary service reference signal.

For (ii), we compared the energy consumption resulting from the proposed method with that from the conventional climate control system used in Pugh Hall – the single-maximum control [26]. We call this control logic the "Pugh Hall control".

TABLE III
ENERGY CONSUMPTION COMPARISON BETWEEN SIMULATIONS WITH AND WITHOUT PROVIDING ANCILLARY SERVICES (AS).

	Without AS	With AS	Increase with AS	
Cooling (kWh)	198.0	198.1	0.1	
Reheat (kWh)	21.7	21.9	0.2	
Fan (kWh)	7.6	7.9	0.3	
Total (kWh)	227.3	228.0	0.7	(0.3 %)

The resulting changes in the daily energy consumption and its breakdown are shown in Table IV. We observe that the proposed approach leads to a slight increase in energy consumption over that of the conventional climate control on the summer day (6%) and a larger increase on the winter day (11%). The reason for this is the model mismatch resulting from computing the cooling load from the supply air flow rate. Time-traces of relevant signals from the simulations are shown in Fig. 10 and Fig. 11. Recall equations (6) to (10), where the calculation of the cooling demand P_{cc} is based on a linearized cooling coil model. The cooling power demand that the scheduler calculates is higher than the true value. This also leads to a higher heating load to maintain indoor temperature of the building. In short, the inaccurate cooling coil model used by the scheduler leads to a high degree of simultaneous heating and cooling. Note that, in the summer simulation, the end-of-day temperature is higher for BYOB-ASC than Pugh Hall control. If the BYOB-ASC achieved a final temperature equals to that of the Pugh Hall control, the energy consumption would be slightly different.

Although it may not be possible to completely eliminate simultaneous cooling and heating with existing HVAC systems, the increase in this inefficient behavior of the proposed approach can be improved by a number of means. These include calibrating different models for different seasons, or through an on-line cooling coil adaptation.

V. Conclusion

Building HVAC systems can provide ancillary services without affecting QoS over a range of timescales due to the large thermal inertia of the building. The proposed scheme is able to provide satisfactory reference tracking without significant effect on the indoor climate for reference signals within the frequency range $f \in [1/(1 \text{ hour}), 1/(10 \text{ minutes})]$. A slight increase in energy consumption is observed compared to the conventional control scheme. Reducing energy consumption through improvement of the control algorithm is an avenue for future work.

The site-specific information required to implement the proposed system in a building is small, and can be obtained from a few weeks of data and system-identification experiments. The system can be deployed easily in a building with a BAS, without requiring installation of any new hardware. One-way communication of the reference signal from the system operator to the building is required, which can be performed, for example, over the public Internet by using technologies such as PJM's Jetstream© [30].

Future work will involve examining the performance of the system in simulations with a more sophisticated model of the chiller than the one used here, and experimental verification. It will also be useful to extend the proposed approach to multi-zone buildings. Another avenue is to extend the algorithm to include energy reduction, so that the proposed system can provide services to both the consumer (energy efficiency) and to the grid (ancillary services), as in [9–11].

TABLE IV
ENERGY CONSUMPTION COMPARISON BETWEEN SIMULATIONS WITH BYOB-ASC (REFERRED TO AS "BYOB") AND PUGH HALL CONTROL (REFERRED TO AS "PUGH").

	Summer day			Winter day		
	Pugh	BYOB	Increase with BYOB	Pugh	BYOB	Increase with BYOB
Cooling (kWh)	198.0	198.0	0.0	56.7	64.8	8.1
Heating (kWh)	8.0	21.7	13.7	123.6	135.0	11.4
Fan (kWh)	7.8	7.6	-0.2	4.4	4.8	0.4
Total (kWh)	213.9	227.3	13.4 (6%)	184.7	204.7	20.0 (11%)

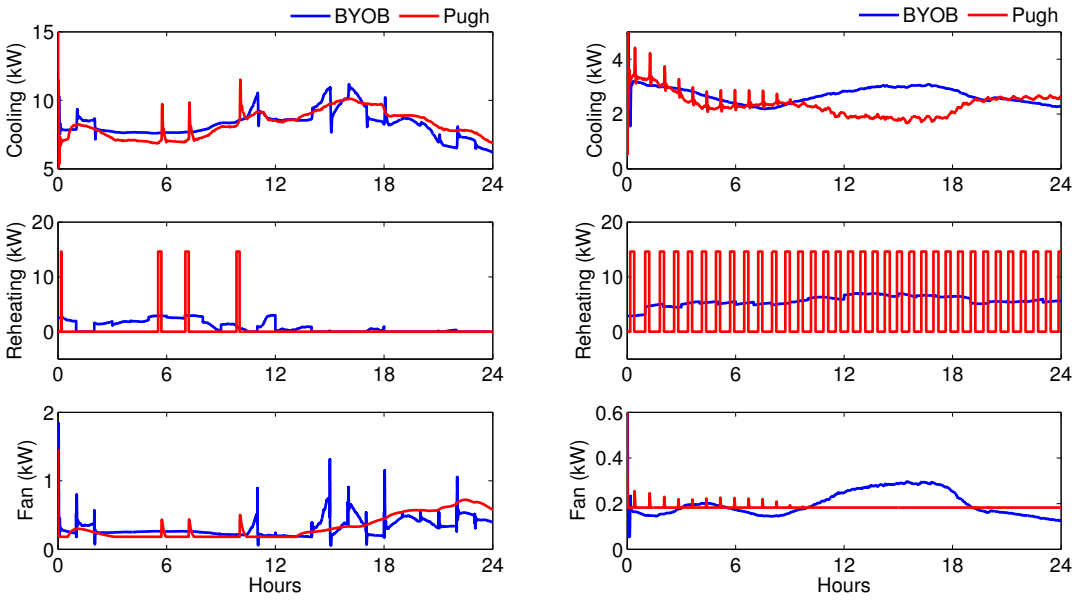

Fig. 10. Power consumption comparison between simulations with BYOB-ASC (referred to as "BYOB") and Pugh Hall control (referred to as "Pugh"). Left: summer simulation; right: winter simulation.

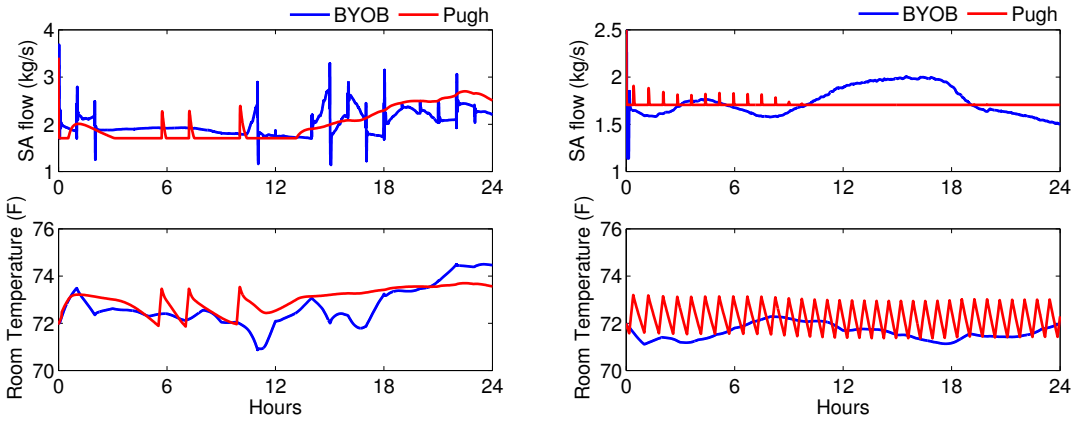

Fig. 11. Temperature and supply air flow rate comparison between simulations with BYOB-ASC (referred to as "BYOB") and Pugh Hall control (referred to as "Pugh"). Left: summer simulation; right: winter simulation. SA refers to supply air.

REFERENCES

[1] J. Mathieu, S. Koch, and D. Callaway, "State estimation and control of electric loads to manage real-time energy imbalance," *IEEE Transactions on Power Systems*, vol. 28, no. 1, pp. 430–440, 2013.

[2] N. Lu, "An evaluation of the HVAC load potential for providing load balancing service," *Smart Grid, IEEE Transactions on*, vol. 3, no. 3, pp. 1263–1270, 2012.

[3] B. Biegel, P. Andersen, T. S. Pedersen, K. M. Nielsen, J. Stoustrup, and L. H. Hansen, "Electricity market optimization of heat pump portfolio," in *Control Applications (CCA), 2013 IEEE International Conference on*. IEEE, 2013, pp. 294–301.

[4] S. Meyn, P. Barooah, A. Busic, Y. Chen, and J. Ehren, "Ancillary service to the grid using intelligent deferrable loads," *Automatic Control, IEEE Transactions on*, vol. 60, no. 11, pp. 2847–2862, 2015.

[5] Y. Kim, L. Norford, and J. Kirtley, "Modeling and analysis of a variable speed heat pump for frequency regulation through direct load control," *Power Systems, IEEE Transactions on*, vol. 30, pp. 397 – 408, 2014.

[6] H. Hao, T. Middelkoop, P. Barooah, and S. Meyn, "How demand response from commercial buildings will provide the regulation needs of the grid," in *50th Annual Allerton Conference on Communication, Control, and Computing (Allerton), 2012*, pp. 1908–1913.

[7] H. Hao, Y. Lin, A. Kowli, P. Barooah, and S. Meyn, "Ancillary service to the grid through control of fans in commercial building HVAC systems," *Smart Grid, IEEE Transactions on*, vol. 5, no. 4, pp. 2066–2074, July 2014.

[8] Y. Lin, P. Barooah, S. Meyn, and T. Middelkoop, "Experimental evaluation of frequency regulation from commercial building HVAC systems," *IEEE Transactions on Smart Grid*, vol. 6, pp. 776 – 783, 2015.

[9] E. Vrettos, F. Oldewurtel, F. Zhu, and G. Andersson, "Robust provision of frequency reserves by office building aggregations," in *IFAC World Congress*, vol. 19, no. 1, 2014, pp. 12 068–12 073.

[10] E. Vrettos, F. Oldewurtel, and G. Andersson, "Robust energy-constrained frequency reserves from aggregations of commercial buildings," *IEEE Transactions on Power Systems*, 2016, DOI:10.1109/TPWRS.2015.2511541.

[11] M. Maasoumy, C. Rosenberg, A. Sangiovanni-Vincentelli, and D. S. Callaway, "Model predictive control approach to online computation of demand-side flexibility of commercial buildings HVAC systems for supply following," in *American Control Conference (ACC), 2014*. IEEE, 2014, pp. 1082–1089.

[12] B. Kirby, "Ancillary services: technical and commercial insights," 2007. [Online]. Available: http://www.science.smith.edu/~jcardell/Courses/EGR325/Readings/Ancillary_Services_Kirby.pdf

[13] Y. Lin, P. Barooah, S. Meyn, and T. Middelkoop, "Demand side frequency regulation from commercial building HVAC systems: an experimental study," in *American Control Conference (ACC), 2015*. IEEE, 2015, pp. 3019–3024.

[14] J. Mathieu, D. Callaway, and S. Kiliccote, "Variability in automated responses of commercial buildings and industrial facilities to dynamic electricity prices," *Energy and Buildings*, vol. 43, pp. 3322–3330, 2011.

[15] S. Borenstein, M. Jaske, and A. Rosenfeld, "Dynamic pricing, advanced metering and demand response in electricity markets," University of California Energy Institute: Center for the Study of Energy Markets, Tech. Rep. CSEMWP-105, 2002.

[16] Y. Lin, P. Barooah, and S. P. Meyn, "Low-frequency power-grid ancillary services from commercial building HVAC systems," in *SmartGridComm, 2013 IEEE International Conference on*. IEEE, 2013, pp. 169–174.

[17] Y. Lin, P. Barooah, and J. Mathieu, "Ancillary services to the grid from commercial buildings through demand scheduling and control," in *American Control Conference (ACC), 2015*. IEEE, 2015, pp. 3007–3012.

[18] P. Barooah, A. Bušić, and S. Meyn, "Spectral decomposition of demand-side flexibility for reliable ancillary services in a smart grid," in *System Sciences (HICSS), 2015 48th Hawaii International Conference on*. IEEE, 2015, pp. 2700–2709.

[19] PJM, "PJM manual 12: Balancing operations, rev. 27," December 2012.

[20] J. Mathieu, P. Price, S. Kiliccote, and M. Piette, "Quantifying changes in building electricity use, with application to demand response," *IEEE Transactions on Smart Grid*, vol. 2, no. 3, pp. 507–518, 2011.

[21] American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc., "ANSI/ASHRAE standard 62.1-2007, ventilation for acceptable air quality," 2007. [Online]. Available: www.ashrae.org

[22] S. Bendapudi, J. Braun, and E. Groll, "A dynamic model of a vapor compression liquid chiller," in *International Refrigeration and Air Conditioning Conference*, 2002.

[23] K. Warwick and D. Rees, *Industrial digital control systems*. Peter Peregrinus Limited, 1988, vol. 37.

[24] I. B. Rhodes, "A tutorial introduction to estimation and filtering," *IEEE Transaction on Automatic Control*, vol. AC-16, no. 6, 1971.

[25] Y. Lin, "Control of commercial building HVAC systems for power grid ancillary service," Ph.D. dissertation, University of Florida, 2014. [Online]. Available: http://www.uflib.ufl.edu/etd.html

[26] American Society of Heating, Refrigerating and Air Conditioning Engineers, "The ASHRAE handbook fundamentals (SI Edition)," 2005.

[27] ——, "ASHRAE Standard 90.1-2004, energy standard for buildings except low rise residential buildings," 2010.

[28] S. Goyal, H. Ingley, and P. Barooah, "Occupancy-based zone climate control for energy efficient buildings: Complexity vs. performance," *Applied Energy*, vol. 106, pp. 209–221, June 2013.

[29] I. Beil, I. Hiskens, and S. Backhaus, "Round-trip efficiency of fast demand response in a large commercial air conditioner," *Energy and Buildings*, vol. 97, pp. 47–55, 2015.

[30] PJM, "Jetstream," http://www.pjm.com/markets-and-operations/etools/jetstream.aspx.

Yashen Lin is a Post-doctoral research fellow in the Department of Electrical Engineering and Computer Science at the University of Michigan, Ann Arbor, MI, USA. He received the B.S. degree in automation from the University of Science and Technology, Beijing, China, in 2009. He received the M.S. and Ph.D. degrees in mechanical engineering from the University of Florida, Gainesville, FL, USA, in 2012 and 2014, respectively. His research interests include power system modeling and optimization, energy storage resources, and building HVAC control.

Prabir Barooah is an Associate Professor of Mechanical and Aerospace Engineering at the University of Florida, where he has been since 2007. He received the Ph.D. degree in Electrical and Computer Engineering in 2007 from the University of California, Santa Barbara. From 1999 to 2002 he was a research engineer at United Technologies Research Center, East Hartford, CT. He received the M. S. degree in Mechanical Engineering from the University of Delaware in 1999 and the B. Tech. degree in Mechanical Engineering from the Indian Institute of Technology, Kanpur, in 1996. Dr. Barooah is the winner of the ASEE-SE (American Society of Engineering Education, South East Section) outstanding researcher award (2012), NSF CAREER award (2010), General Chairs' Recognition Award for Interactive papers at the 48th IEEE Conference on Decision and Control (2009), best paper award at the 2nd Int. Conf. on Intelligent Sensing and Information Processing (2005), and NASA group achievement award (2003).

Johanna L. Mathieu (S'10M'12) received the B.S. degree in ocean engineering from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2004 and the M.S. and Ph.D. degrees in mechanical engineering from the University of California, Berkeley, USA, in 2008 and 2012, respectively. She is an Assistant Professor in the Department of Electrical Engineering and Computer Science at the University of Michigan, Ann Arbor, MI, USA. Prior to joining the University of Michigan, she was a postdoctoral researcher at the Swiss Federal Institute of Technology (ETH) Zurich, Switzerland. Her research interests include modeling, estimation, control, and optimization of demand response and energy storage resources.