Ancillary Services to the Grid from Commercial Buildings through Demand Scheduling and Control

Yashen Lin, Prabir Barooah, and Johanna L. Mathieu

***Abstract*— How can a building Heating, Ventilation, Air Conditioning (HVAC) system vary its real time power consumption to provide ancillary services to the power grid without sacrificing occupant comfort? Prior work showed how this can be done if the reference power variation is of high frequency (seconds to a few minutes) so that the climate control system filters out the disturbance. 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. We propose a receding horizon approach to schedule the baseline cooling and heating power of a building based on weather forecasts. A lower level controller is then used to track the scheduled baseline plus ancillary service reference signal. Periodic updates to the scheduler based on measurements ensure quality of service in spite of forecasting errors. The algorithm is tested in simulation. Results show that ancillary service 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.**

I. Introduction

Ancillary services are needed to correct the mismatch between demand and supply in a power grid to ensure the functionality and reliability of the grid. Integrating a large amount of volatile renewables into the power grid will require a larger amount of ancillary services to handle this volatility [1, 2]. Traditionally, ancillary services are provided by fast ramping generators. An alternative is to explore demand-side flexibility [3, 4] which may have less environmental impact and cost in the long run. Interest in demand side resources providing ancillary services is long standing. Florida Power & Light's On Call® program is one of the early attempts at using demand side resources [5]. Recent research on providing various forms of ancillary services with on-off loads include [6–8], which consider thermostatic loads such as residential air conditioners, heat pumps, refrigerators etc., and [9] which considers deferrable loads, such as pool pumps, with local intelligence.

This paper builds on prior work on using commercial building HVAC systems for providing ancillary services [10–12]. Commercial buildings account for about 40% of the total electricity consumption in the U.S. [13]. The power consumption in Variable Air Volume (VAV) Heating, Ventilation, Air Conditioning (HVAC) systems, which serve 30% of all commercial building floor space in the U.S. [14], can be varied continuously between a low and high value using variable speed drives. This feature makes them particularly well-suited for sophisticated control. Moreover, many commercial buildings are equipped with Building Automation Systems (BAS), making the task of implementing additional control algorithms easy and inexpensive. Finally, commercial buildings have high thermal inertia, which can be translated to effective energy storage much like a large battery.

In commercial building HVAC systems, multiple pieces of equipment can be used to provide ancillary services at different timescales. The supply air fan has fast dynamics, and is suitable for high frequency ancillary services; see [10, 15]. Heat pumps with variable speed drives are another potential resource [16]. Chillers, even those without variable speed drives, can be used to provide ancillary services by indirectly varying the load on them [11]. Due to their slow dynamics, chillers are useful for providing service at slow time scales.

When the reference signal is high-frequency compared to the thermal dynamics of the building and its climate control system, a low pass filter can be used to estimate the baseline power [10]. However, when the reference signal frequency overlaps with the bandwidth of the building, the baseline estimation problem is more challenging. Building climate control systems are designed to react to disturbances on the time scale of a few minutes or longer. Estimating baselines on these timescales requires prediction, typically with statistical baseline models parameterized with historical building data [17, 18]. Errors in predictions make it hard to separate the baseline and control response [19].

In this paper, we propose a novel approach to the baseline estimation problem and a new algorithm to effectively control commercial building HVAC systems to provide ancillary services on timescales of a few minutes to an hour. Instead of estimating the baseline, the baseline is *scheduled* ahead of time based on weather forecasts. Note that Borenstein et al. [20] proposed a similar concept for demand response financial settlement purposes called Build-Your-Own (BYO) baseline. After scheduling the baseline, a lower level Power Tracking Controller (PTC) is used for tracking the scheduled baseline plus the filtered reference signal. The scheduler updates the baseline periodically based on indoor climate measurements to ensure indoor climate conditions remain comfortable. This also ensures robustness to load forecast error and other uncertainties. The benefits of the proposed method are two-fold: (i) the baseline is clearly defined which makes implementing and evaluating the service provided possible, and (ii) by smart scheduling, we can minimize the energy consumption of the HVAC system. Following Borenstein et al., we call our approach a Bring-Your-Own-

YL and PB are with the Department of Mechanical and Aerospace Engineering, University of Florida; yashenl@umich.edu, pbarooah@ufl.edu. JM is with the Department of Electrical Engineering and Computer Science, University of Michigan; jlmath@umich.edu. This research was supported by the NSF grant ECCS-0925534.

Baseline (BYOB) scheme.

The algorithm is tested in simulation where the model is built to resemble a portion of Pugh Hall on the University of Florida campus. Results show the algorithm could provide satisfactory ancillary services in the frequency range of $f \in [1/(1 \text{ hour}), 1/(10 \text{ minutes})]$ while maintaining comfortable indoor climate.

The rest of the paper is organized as follows: Section II describes the proposed algorithm; Section III present a stability analysis of the algorithm; Section IV introduces the model used in the simulation, the simulation setup, and the results; Section V concludes the paper and discusses future work.

II. Control Architecture

We start with a description of VAV HVAC systems. The algorithm leaves certain parts of the climate control system untouched while overriding others.

A. VAV HVAC systems

In a variable air volume HVAC system, the indoor climate and Indoor Air Quality (IAQ) are maintained by varying the flow rate of air through the building. Fresh air brought in from the outside and return air collected from the zones are mixed and sent through a cooling coil in the Air Handling Unit (AHU), where the air is cooled and dehumidified. There may also be a reheat coil to increase the temperature after it is dehumidified. The conditioned air is then sent to the terminal VAV boxes by a supply air fan for distribution to zones. We consider a chilled-water based HVAC system, where a chiller produces chilled water which is then used in the AHU cooling coils to condition the air. Similarly, a boiler produces the hot water used for reheating.

Among the three main sources of power consumption in an HVAC system – mechanical, cooling, and heating – we only consider chiller power (cooling) for providing ancillary services. Cooling and heating power dominate the mechanical power, and in heating is often provided by steam whose generation uses little electricity.

B. Proposed BYOB system

As in [10–12], it is assumed that a reference signal for power deviation, $\delta P^{BA}(t)$, is transmitted by the Balancing Authority (BA) to all the ancillary services providers, including smart buildings. For demand side resources, this signal is the desired deviation of the loads' power consumption from its baseline value, i.e., what the loads would have consumed if they were not providing ancillary services. We assume that $\delta P^{BA}(t)$ is band-pass filtered at every building to obtain a reference signal $\delta P^r(t)$ (in Watts) with magnitude and frequency that is appropriate for that building. The objective of the BYOB control system is twofold: to vary the power consumption of the HVAC system so that the deviation from its baseline tracks δP^r, and ensure that the indoor climate and IAQ are maintained within their bounds.

As discussed in Section II-A, we only consider chiller power for ancillary services in this paper. The total chiller

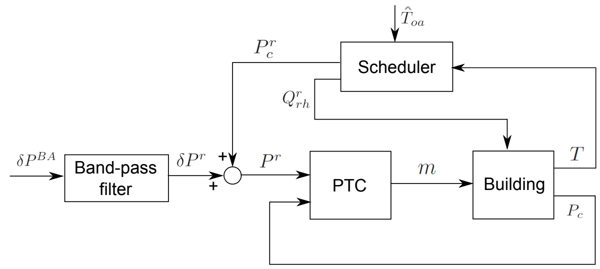

Fig. 1. Schematic illustration of the proposed algorithm.

power $P_c(t)$ is the sum of the baseline power $P_c^b(t)$ and the power deviation $\delta P(t)$ introduced by the BYOB system for providing ancillary service, i.e.,

$$P_c(t) = P_c^b(t) + \delta P(t) \tag{1}$$

As described in Section I, the baseline is *scheduled* ahead of time to match the forecasted thermal load. The scheduled baseline chiller power is called P_c^r. The reference command for the total power consumption is the sum of the scheduled baseline and the ancillary services reference signal:

$$P^r(t) = P_c^r(t) + \delta P^r(t). \tag{2}$$

The goal of the proposed system, whose schematic is shown in Fig. 1, is to ensure that the output P_c tracks P^r, while maintaining indoor climate. A low level controller, the PTC, varies the airflow rate so that the power consumption of the HVAC system P_c tracks the reference P^r. A high level "scheduler" schedules references for cooling and reheating power consumption, P_c^r and Q_{rh}^r (both have units of Watts) ahead of time. Note that to ensure a comfortable room climate, the scheduler computes reference signals for both chiller power and reheat power. The scheduled reheat power will be provided by local controllers in the BAS. Since reheat power is not used for providing ancillary services, we will focus on designing the PTC for chiller power tracking.

The existing building climate control system that is used to vary the air flow rate to maintain indoor climate is overridden for power tracking. In order to maintain the indoor climate, the scheduled power consumption is periodically updated using receding-horizon optimization based on load forecasts and corrections. The steps are described next. Time is divided into a number of slots for scheduling and implementation purposes, with length Δt_S and Δt_I, respectively.

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-C.1.
2) Decide the desired baseline power during $\mathcal{T}_k^S$ by solving an optimization problem: minimize cooling and reheating energy during $\mathcal{T}_k^S$ while ensuring that thermal comfort and IAQ constraints are satisfied and actuation limits are respected.
3) During the implementation period $\mathcal{T}_k^I := (k\Delta t_I,\ (k+1)\Delta t_I]$, track $P^r(t)$ by varying the supply air flow rate.

The PTC is designed for this purpose.

4) At the end of implementation period, update the load estimation for the next scheduling period $\mathcal{T}_{k+1}^S$, and go back to step 1.

C. Baseline scheduling

The primary task for an HVAC system is to maintain a comfortable indoor climate. Thus, forecasting the load is a key part of scheduling the baseline.

1) Load forecasting model: We use the baseline model developed in [21] to forecast the future load on the HVAC system based on outdoor air temperature forecasts and time-of-week, which acts as a proxy for occupant behavior and appliance use. In [19, 21], the model is used to predict the baseline power consumption of *a whole building* during a demand response event that occurred *in the past*, using past outdoor air temperature measurements and time-of-week. In contrast, here we use it to forecast the *future thermal load* of the building.

Before describing the model, we have to define the term "thermal load" precisely and describe how to measure it so that measurements can be used to fit a model. 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:

$$C\dot{T}(t) = -\,Q_{cc}(t) + Q_{rh}(t) + Q_l(t) \tag{3}$$

where $Q_{rh}(t), Q_{cc}(t)$ are the rate of heat provided to and extracted from the building, respectively, by the heating and cooling coils, and the last term $Q_l(t)$ is the rate of heat entering the building from all other sources, such as from outside air, solar irradiation, occupants and plug loads.

Ignoring the modeling error, $Q_l(t)$ is called the *thermal load* experienced by the HVAC system. It is the rate of heat the HVAC system has to remove from the building in order to maintain indoor temperature at a constant set-point, T^{sp}, so that $\dot{T}(t) = 0$. Thus, assuming that the existing climate controller is able to do its job perfectly, the following relationship should hold:

$$0 = -Q_{cc}(t) + Q_{rh}(t) + Q_l(t) \tag{4}$$

This equation gives us $Q_l(t)$ from measurements of $Q_{cc}(t)$ and $Q_{rh}(t)$, which are obtained from energy meters installed in the cooling and heating systems.

Once the question of obtaining measurements for the building's load is resolved, the next question is fitting a model to this data. The model in [21] is a linear regression model, where the load is a piecewise linear function of outdoor air temperature and of time-of-week. Ordinary least squares 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 outdoor air temperature.

2) Baseline scheduling at $t = 0$: The baseline power consumption is scheduled by solving an optimization problem: minimize energy consumption during the k-th scheduling time interval $\mathcal{T}_k^S$ while satisfying constraints such as thermal comfort, ventilation requirements, and equipment saturation limits. 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 power consumption we consider in the optimization is the sum of cooling coil power and reheat power, i.e., $Q_{cc}(t) + Q_{rh}(t)$. The reheat power is assumed to be controlled directly, which makes $Q_{rh}(t)$ a decision variable. The cooling coil power is controlled indirectly by controlling the mass flow rate of air through the cooling coil.

The problem of minimizing the energy consumption can be stated as follows, for $t \in \mathcal{T}_0^S = [0, \Delta t_S]$:

$$\underset{m(t), Q_{rh}(t)}{\text{minimize}} \int_{\mathcal{T}_0^S} (Q_{cc}(t) + Q_{rh}(t))dt \tag{5}$$

subject to

$$-Q_{cc}(t) + Q_{rh}(t) + \hat{Q}_l(t) = 0 \tag{6}$$

$$m(t) \in [m_{lb}, m_{ub}], \quad Q_{rh}(t) \in [Q_{lb}, Q_{ub}] \tag{7}$$

where $m(t)$ is the supply air mass flow rate; m_{lb} and m_{ub} are the bounds for supply air flow rate; Q_{lb} and Q_{ub} are the bounds for reheat power; $\hat{Q}_l(t)$ is the forecast of load from the load forecasting model described in Section II-C.1. The cooling power $Q_{cc}(t)$ equals the product of $m(t)$ and enthalpy difference between the air stream before the cooling coil (*mixed air*) and after the cooling coil (*discharge air*):

$$Q_{cc}(t) = m(t)(h_{ma}(t) - h_{da}(t)) \tag{8}$$

where h_{ma} and h_{da} are the enthalpies of the mixed air and discharged air. Under the assumptions that all set points are constants and local control loops maintain their outputs at their set points, the specific enthalpies do not depend on the actuation $m(t)$. The energy minimization is then 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 $t \in \mathcal{T}_0^S$:

$$m^r(t), Q_{rh}^r(t) = \underset{m(t), Q_{rh}(t)}{\arg\min} \ (Q_{cc}(t) + Q_{rh}(t)) \tag{9}$$

subject to the load and actuator constraints specified by (6) and (7). This is a linear program in two decision variables, $m(t)$ and $Q_{rh}(t)$, and can be solved easily. The scheduled baseline cooling power $Q_{cc}^r(t)$ is obtained from (8) once $m^r(t)$ is determined. The chiller dynamics are modeled as a first order LTI system, with Q_{cc} as input and P_c as output. With this model, the scheduled chiller power $P_c^r(t)$ can be calculated from $Q_{cc}^r(t)$.

D. Baseline update

In practice, the scheduled power will not lead to an indoor temperature exactly equal to the set-point due to uncertainties such as load forecast error. In subsequent scheduling periods, these uncertainties are accounted for by examining how far the actual space temperature varies from the set-point, and adding correction terms.

The baseline is updated at any scheduling period $\mathcal{T}_k^S$, $k \geq 1$, by solving the optimization problem (9), with only one difference, specifically the load constraint (6) is replaced by:

$$-Q_{cc}(t) + Q_{rh}(t) + \hat{Q}_l(t) + Q_{c1}^k(t) + Q_{c2}^k(t) = 0, \quad t \in \mathcal{T}_k^S \tag{10}$$

where $Q_{c1}^k(t), Q_{c2}^k(t)$ are two correction terms. We now describe these correction terms and the rationale behind their design.

The uncertainties in the temperature dynamics (3) are captured by a disturbance term, $\delta Q_l(t)$, so that (3) now becomes:

$$C\dot{T}(t) = -Q_{cc}^r(t) + Q_{rh}^r(t) + \hat{Q}_l(t) + \delta Q_l(t) \tag{11}$$

where the term $\delta Q_l(t)$ also captures any deviation of the actual cooling and heating power from their scheduled values. Since the scheduled $Q_{cc}^r(t)$ and $Q_{rh}^r(t)$ have to satisfy the constraint (10), we have:

$$C\dot{T}(t) = -Q_{c1}^k(t) - Q_{c2}^k(t) + \delta Q_l(t) \tag{12}$$

Note that for the first scheduling period, the correction terms are 0.

The first correction term is chosen as

$$Q_{c1}^k(t) := \gamma_1 \dot{T}(k\Delta t_I) e^{-\lambda(t - k\Delta t_I)}, \quad t > k\Delta t_I \tag{13}$$

where γ_1 and λ are design parameters. The rationale behind this term is that $\gamma_1 \dot{T}$ estimates load forecast error. We expect the actual load tends to return to the estimated value, i.e., $\delta Q_l(t)$ tends to go to 0 for large t.

The second correction term is designed to remove the total extra heat that has remained in the building at the end of the the last (k-1 -th) implementation period, which is simply

$$E_{end} = C(T(k\Delta t_I) - T^{sp})$$

The correction term is chosen as

$$Q_{c2}^k(t) := \begin{cases} \gamma_2 E_{end} & t \in (k\Delta t_I,\ k\Delta t_I + \frac{1}{\gamma_2}] \\ 0 & t \in (k\Delta t_I + \frac{1}{\gamma_2}, k\Delta t_I + \Delta t_S] \end{cases} \tag{14}$$

where γ_2 is a design parameter which indicates how fast we would like the temperature to be driven back to the set-point. If we wish to return the zone temperature to its set-point quicker, a larger γ_2 should be used.

E. PTC design

The objective of the PTC is to make the chiller power track the total reference signal $P^r(t)$ during every implementation period: $t \in \mathcal{T}_k^I$. A dynamic compensator is used for this purpose, with power tracking error $P^r - P_c$ as input and air flow rate command m^r as output. Classical compensator design techniques are used to design the PTC based on a linear plant model, which is obtained by linearizing a non-linear model of the HVAC system around a nominal operating point.

III. Stability analysis

The baseline update provides a feedback mechanism. Consider the case where there is no update: the load $\hat{Q}_l(t)$ is pre-calculated from the load model based on weather forecasts. The reference signal is then calculated from (9). This control strategy works in a feed-forward fashion, where the input ($Q_{cc}^r(t)$ and $Q_{rh}^r(t)$) is decided ahead of time and does not depend on the real time room temperature measurement.

Now consider the case with update. The dynamics after adding the two correcting terms are given in (12). For convenience of analysis, we let $x_1(t) = T(t) - T^{sp}$, 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 the definition in (13) and (14), we have:

$$Q_{c1}^k(t) = \gamma_1 x_2(k\Delta t_I) e^{-\lambda(t - k\Delta t_I)} \tag{15}$$

$$Q_{c2}^k(t) = \gamma_2 C x_1(k\Delta t_I) \tag{16}$$

We assume $\frac{1}{\gamma_2} \geq \Delta t_I$ so that the second correction term $Q_{c2}^k(t)$ remains constant during $\mathcal{T}_k^I$. This simplifies the expression. In the case of $\frac{1}{\gamma_2} < \Delta t_I$, $Q_{c2}^k(t)$ becomes a step function which drops to 0 at $t = k\Delta t_I + \frac{1}{\gamma_2}$. From (12),

$$x_2(t) = \dot{T}(t) = \frac{1}{C}(-Q_{c1}^k(t) - Q_{c2}^k(t) + \delta Q_l(t)) \tag{17}$$

We now define the discrete states:

$$x_1^k = x_1(k\Delta t_I) \qquad x_2^k = x_2(k\Delta t_I). \tag{18}$$

At time $t = (k+1)\Delta t_I$, $x_1^{k+1} = x_1^k + \int_{k\Delta t_I}^{(k+1)\Delta t_I} x_2(t)dt$. By using (17), it follows that

$$x_1^{k+1} = (1 - \gamma_2 \Delta t_I)x_1^k + \frac{\gamma_1}{C\lambda}(e^{-\lambda \Delta t_I} - 1)x_2^k + \delta W^{k+1} \tag{19}$$

where

$$\delta W^{k+1} := \frac{1}{C}\int_{k\Delta t_I}^{(k+1)\Delta t_I} \delta Q_l(t)dt \tag{20}$$

Also, we have:

$$\begin{aligned} x_2^{k+1} &= \dot{T}((k+1)\Delta t_I) \\ &= \frac{1}{C}(-Q_{c1}^k((k+1)\Delta t_I) - Q_{c2}^k((k+1)\Delta t_I) \\ &\quad + \delta Q_l((k+1)\Delta t_I)) \\ &= -\gamma_2 x_1^k - \frac{\gamma_1}{C} e^{-\lambda \Delta t_I} x_2^k + \frac{1}{C}\delta Q_l^{k+1} \end{aligned} \tag{21}$$

Combining, we get

$$\begin{bmatrix} x_1^{k+1} \\ x_2^{k+1} \end{bmatrix} = \begin{bmatrix} 1 - \gamma_2 \Delta t_I & \frac{\gamma_1}{C\lambda}(e^{-\lambda \Delta t_I} - 1) \\ -\gamma_2 & \frac{\gamma_1}{C} e^{-\lambda \Delta t_I} \end{bmatrix} \begin{bmatrix} x_1^k \\ x_2^k \end{bmatrix} + \begin{bmatrix} \delta W^{k+1} \\ \frac{1}{C}\delta Q_l^{k+1} \end{bmatrix} \tag{22}$$

The values of γ_1, γ_2, and λ used in the simulation studies reported in this paper are given in Table I. For those values of the parameters, the eigenvalues of the state matrix in (22) are 0.52 and 0.04, which shows the discrete system is BIBO stable.

TABLE I

PARAMETER VALUES USED IN SIMULATION.

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	T^{sp}	72 ^oF

IV. SIMULATION STUDY

The algorithm is studied in simulations on a high-fidelity model of a building and its HVAC system.

A. Reference signals

In this simulation, the Area Control Error (ACE) signal from PJM on 05/04/2009 is used as δP^{BA}, which is bandpass filtered by a 10^{th}-order Butterworth filter with a pass band of $f \in [1/(1 \text{ hour}), 1/(10 \text{ minutes})]$ and a passband gain of 5.5×10^{-6}. The resulting δP^r has a maximum magnitude of 5kW.

B. Building and HVAC system model for simulations

The model of the building and its HVAC system used in simulations is constructed and calibrated to resemble a large zone in Pugh Hall at the University of Florida that is serviced by a dedicated AHU. The model consists of the dynamics of the zone, supply air fan, distribution duct, cooling coil, and chiller, as well as the feedback interconnections of the components. Due to lack of space, we do not describe the model in detail here; the time constant of the chiller dynamics is chosen to be 200 seconds, and other components of the model are described in [11].

C. Controller

Recall that the algorithm has two main components, a baseline scheduler and a PTC. The baseline scheduler needs the data-driven load model described earlier. Historical data of Q_{cc} and Q_{rh} are collected from sensors installed in Pugh Hall. Weather data are collected from *www.wunderground.com*. Two weeks of data (01/13/2014 to 01/26/2014) are used in estimating the parameters of the model. Though little data are used to parameterize the model, the receding horizon approach provides robustness to model inaccuracy. Values of other parameters needed by the baseline scheduler and updater are shown in Table I.

The model of the HVAC system described in Section IV-B is linearized to obtain the transfer function $H_{m,p}(s)$ (m_a^d to P_c) for controller design. The controller C_{PTC} is designed based on $H_{m,p}(s)$, so that the sensitivity function of the closed loop (P^r to P_c) is $S(s) = \frac{s}{s+0.1}$, which makes the tracking error less than -3 dB when $\omega < 0.1$ rad/s. Although designed with a linearized model, the PTC is applied to the high fidelity non-linear model of the HVAC system in the simulations.

D. Performance metrics

The dual goals of the algorithm are frequency regulation to the grid, while maintaining IAQ. Metrics to quantify performance are described here. Let $\delta P := P_c - P^r$ be the measured chiller power deviation from the baseline. A natural metric to quantify the quality of tracking error $e(t) := \delta P^r(t) - \delta P(t)$ is the ratio

$$r_R = \frac{1}{\max |\delta P^r|} \sqrt{\int_0^\tau e(t)^2 dt}. \tag{23}$$

We also evaluate the algorithm using PJM Interconnection's performance score, based on the formula given in [22]. The total performance score S_t is the mean of three scores: correlation score S_c, delay score S_d, and precision score S_p. A score of $S_t \geq 0.75$ is required to pass the PJM test.

For indoor climate, we use the temperature violation D_T defined in [23]. This score is based on the minimum and maximum temperature allowed when the building is occupied, set to $70°F$ and $75°F$ according to the thermal comfort specifications described in [24].

Variation in supply air flow rate is quantified by

$$\delta m_{avg} = \frac{1}{\tau} \int_0^\tau \left| \frac{m(t) - m^r(t)}{m^r(t)} \right| dt \tag{24}$$

where m, m^r denote the actual and scheduled (reference) supply air flow rates.

E. Results

Tracking performance is shown in Fig. 2. Performance metrics are shown in Table II. We see that our controller has a performance score of 0.961; cf. PJM's requirement of 0.75.

The resulting room temperature and airflow rate are shown in Fig. 3. The room temperature remained in the comfortable range during the whole simulation, so the temperature violation D_T is 0. The average variation in supply air flow rate is around 15% of the scheduled baseline.

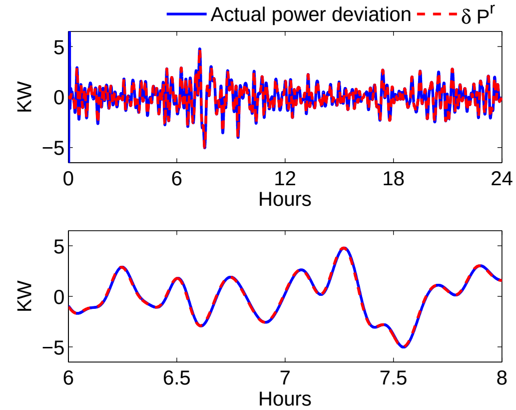

Fig. 2. Power deviation tracking performance. The 2nd plot is a 2-hour close-up.

TABLE II

PERFORMANCE RESULTS FOR SIMULATIONS.

δP^r	r_R	PJM performance score			
		S_c	S_d	S_p	S_t
Filtered ACE	0.016	0.999	0.953	0.931	0.961

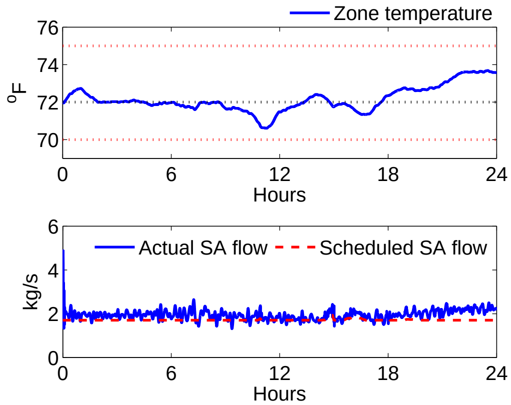

Fig. 3. Room temperature and supply air (SA) flow rate during the simulation. The red horizontal lines in top figure indicate the comfortable temperature range. The black horizontal line in top figure indicates the temperature set-point.

V. Conclusion and Future Work

The proposed approach avoids the baseline estimation problem by scheduling the baseline and periodically updating it based on measurements in a receding horizon optimization framework. Load is forecasted from a data-driven regression model, which is calibrated from cooling and heating energy use data. Load forecasting with the model only requires ambient temperature forecasts, which are available from weather forecasts. Although there are many sources of uncertainty in the load forecast, the feedback introduced through periodic baseline update makes the overall scheme robust to these uncertainties. In the future, we plan to develop methods for on-line estimation of the baseline instead of scheduling it, and compare the two architectures.

The scheduling process requires very little information about the building - only the total thermal capacitance of the building is needed which can be estimated from data. This feature makes the algorithm easily deployable in any building since the calibration effort required to determine a building specific model is small. The algorithm can be deployed easily in a building with a BAS, without requiring installation of any new hardware. The reference signal from the BA to the building can be communicated over the existing public Internet.

In this paper, our algorithm was tested in simulations; experimental investigation are planned in the future. In addition, in the case of an off-site chiller, a transport delay occurs between actuation (air flow rate variation) and power consumption. Addressing the reference tracking problem in presence of this delay is another research avenue.

References

[1] Y. Makarov, C. Loutan, J. Ma, and P. De Mello, "Operational impacts of wind generation on California power systems," *IEEE Transactions on Power Systems*, vol. 24, no. 2, pp. 1039–1050, 2009.

[2] D. Halamay, T. Brekken, A. Simmons, and S. McArthur, "Reserve requirement impacts of large-scale integration of wind, solar, and ocean wave power generation," *IEEE Transactions on Sustainable Energy*, vol. 2, no. 3, pp. 321–328, Jul. 2011.

[3] P. Barooah, A. Bušić, and S. Meyn, "Spectral decomposition of demand side flexibility for reliable ancillary service in a smart grid," in *48th Hawaii International Conference on Systems Science*, January 2015.

[4] D. Callaway and I. Hiskens, "Achieving controllability of electric loads," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 184–199, 2011.

[5] J. R. Stitt, "Implementation of a large-scale direct load control system-some critical factors," *Power Apparatus and Systems, IEEE Transactions on*, vol. PAS-104, no. 7, pp. 1663–1669, July 1985.

[6] 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.

[7] 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.

[8] 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.

[9] S. Meyn, P. Barooah, A. Bušić, Y. Chen, and J. Ehren, "Ancillary service to the grid from intelligent deferrable loads," 2015, in press.

[10] 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.

[11] 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.

[12] 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.

[13] US Department of Energy, "Buildings energy data book," http://buildingsdatabook.eren.doe.gov/, 2011.

[14] "Commercial buildings energy consumption survey (CBECS): Overview of commercial buildings, 2003," Energy information administration, Department of Energy, U.S. Govt., Tech. Rep., December 2008. [Online]. Available: http://www.eia.doe.gov/emeu/cbecs/cbecs2003/overview1.html

[15] 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*, 2015, accepted.

[16] 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*, 2014.

[17] K. Coughlin, M. Piette, C. Goldman, and S. Kiliccote, "Statistical analysis of baseline load models for non-residential buildings," *Energy and Buildings*, vol. 41, no. 4, pp. 374–381, Apr. 2009.

[18] J. Granderson and P. Price, "Evaluation of the predictive accuracy of five whole-building baseline models," Lawrence Berkeley National Laboratory, Berkeley, CA, Tech. Rep. LBNL-5886E, August 2012.

[19] 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.

[20] 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.

[21] 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.

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

[23] 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.

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