Ancillary Service to the Grid through Control of Fans in Commercial Building HVAC Systems

He Hao, Yashen Lin, Anupama Kowli, Prabir Barooah, and Sean Meyn

Abstract—The thermal storage potential in commercial buildings is an enormous resource for providing various ancillary services to the grid. In this paper, we show how fans in Heating, Ventilation, and Air Conditioning (HVAC) systems of commercial buildings alone can provide substantial frequency regulation service, with little change in their indoor environments. A feedforward architecture is proposed to control the fan power consumption to track a regulation signal. The proposed control scheme is then tested through simulations based on a calibrated high fidelity non-linear model of a building. Model parameters are identified from data collected in Pugh Hall, a commercial building located on the University of Florida campus. For the HVAC system under consideration, numerical experiments demonstrate how up to 15% of the rated fan power can be deployed for regulation purpose while having little effect on the building indoor temperature. The regulation signal that can be successfully tracked is constrained in the frequency band $[1/\tau_0, 1/\tau_1]$, where $\tau_0 \approx 3$ minutes and $\tau_1 \approx 8$ seconds. Our results indicate that fans in existing commercial buildings in the U.S. can provide about 70% of the current national regulation reserve requirements in the aforementioned frequency band. A unique advantage of the proposed control scheme is that assessing the value of the ancillary service provided is trivial, which is in stark contrast to many demand-response programs.

Index Terms—Ancillary Service, frequency regulation, commercial buildings, HVAC system, demand response.

I. INTRODUCTION

The electric grid will be subject to more and more volatility from the introduction of renewable energy resources. Hence reliability of the grid will require more ancillary services through generation, as well as flexible consumption. These statements are heard in energy conferences around the world, and in the introductory paragraph of numerous papers. What is often not realized is that each person in the audience, or each author of those papers, is sitting in a vast energy storage device - a building. The thermal storage in buildings is an enormous untapped resource for providing ancillary services to the grid. Moreover, buildings account for 75% of total electricity consumption in the United States [1]. Buildings are, therefore, a natural candidate for demand-side flexibility.



Fig. 1. ACE and total deployed regulation signal within ERCOT (Electric Reliability Council of Texas) [3].

The proper functioning of a power grid requires continuous matching of supply and demand, in spite of the randomness of electric loads and the uncertainty of variable generations. Electric ancillary services, such as frequency regulation and load following, play a crucial role in managing the supply-demand balance. In particular, frequency regulation is deployed on the fastest time-scale (seconds to one minute) to correct the short-term fluctuations in load and generation to maintain system frequency within prescribed limits [2]. This service has been traditionally provided by relatively fast-responding generators by tracking a regulation signal sent by the grid operator. The regulation signal is constructed from the *area control error* (ACE) which measures the amount of (positive or negative) MWs needed in the power system.

In practice, most regulation dispatching algorithms intentionally damp the rapidly moving ACE to accommodate the ramping constraints of the participating generators and to prevent wear and tear of their mechanical equipments. Fig. 1 shows a typical ACE pattern, along with the deployed regulation signal in ERCOT. As shown in the figure, the ACE fluctuates up and down much more frequently and its magnitude is smaller than that of the deployed regulation signal. The reason for the difference is mainly because traditional generators have slow ramping rates; they cannot track the fast changing ACE signal very well, which results in higher frequency regulation procurements [3]. This issue has been recognized in the power community. The FERC (Federal Energy Regulatory Commission) order 755 has been issued to recommend the system operators to pay more for faster regulation resources ("pay for performance").

Studies indicate that increased reliance on renewable generation introduces greater volatility and uncertainty in power system dynamics and imposes additional regulation requirements on the grid [4]–[6]. However, the regulation requirements

He Hao was with the Department of Mechanical and Aerospace Engineering, University of Florida, and is now with Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA 94720, USA, email: hehao@berkeley.edu. Yashen Lin and Prabir Barooah are with Department of Mechanical and Aerospace Engineering, University of Florida, Gainesville, FL 32611, USA. Anupama Kowli is with Department of Electrical Engineering, Indian Institute of Technology Bombay, Powai, Mumbai, India. Sean Meyn is with Department of Electrical and Computer Engineering, University of Florida, Gainesville, FL 32611, USA. This work was supported by the National Science Foundation through Grants CNS-0931885 and ECCS-0925534.

can be lowered if faster responding resources are available [7]. It was shown if CAISO (California independent system operator) dispatched fast responding regulation resources, it could reduce its regulation procurement by as much as 40% [8]. These factors coupled with the search for cleaner sources of flexibility as well as regulatory developments such as FERC order 755 have garnered a growing interest in tapping the fast response potential demand-side resources. In this paper, we argue that (a) commercial building HVAC systems can be manipulated for regulation service on faster timescales more effectively than generators, and, (b) commercial buildings can provide this service at a very low cost.

Although providing ancillary service by managing loads of *residential buildings* has received a lot of attention (see [9]–[12] and references therein), the literature on provision of ancillary service by controlling *commercial buildings* is meager. Many load control mechanisms implemented in utilities and explored in the literature are primarily concerned with low frequency changes in demand such as peak load shaving [13], [14]. Various open-loop load management strategies to reduce commercial building energy consumption in response to grid requirements have been reported in [15], [16].

In this paper, we focus on high frequency (seconds to minutes) load changes in commercial buildings, so as to provide regulation service to the grid. The choice of commercial buildings is motivated by several factors. First, a commercial building can provide a larger potential of ancillary service (compared to a residential building) due to its larger thermal inertia and larger load. Second, approximately one third of the commercial building floor space is equipped with variable frequency drives that operate the HVAC equipment. Their speeds and power can be varied quickly and continuously, instead of in an on/off manner. This is a crucial advantage for providing regulation services with accurate tracking, since the regulation signal to be tracked changes in the order of seconds. Third, a large fraction of commercial buildings in the United States are equipped with Building Automation Systems. These systems can keep track of real-time HVAC system power consumption, and manipulate the control variables needed for providing regulation services, without requiring additional equipment such as smart meters. Ancillary services can thus be provided at very low cost; these are obtained as a simple add-on to the current HVAC control system.

In this paper we showcase the feasibility of extracting regulation service from commercial buildings. We consider power consumption of the supply fans in the HVAC system as the only source of flexibility. A feedforward control architecture is proposed, wherein the fan speed commanded by the building's existing control system is modified so that the change in the fan's power consumption - both decrease and increase - tracks the regulation signal. A simplified dynamic model of a building's HVAC system is used for control design. The model parameters are identified from data collected from a commercial building on the University of Florida campus (Pugh Hall). The controller is then tested on a high fidelity non-linear model constructed from the same building. The results show that the simplified model is adequate for the purpose of control; the controller performs on the complex



Fig. 2. Schematic of a typical commercial building HVAC system that services 11 zones.

model as predicted by the simplified model.

A main contribution of the paper is a characterization of the inherent limitations of providing frequency regulation by commercial buildings with the proposed method, which can be quantified in terms of bandwidth of regulation. For instance, the proposed demand response mechanism through fan speed variation occupies a higher range in the frequency axis, since it is executed through automatic control with high bandwidth actuators. To ensure the comfort of occupants, and to manage stress on HVAC equipment, both upper and lower bounds on bandwidth are necessary. Based on simulation experiments, this bandwidth is estimated to be $[1/\tau_0, 1/\tau_1]$, where $\tau_0 \approx$ 3 minutes, and $\tau_1 \approx 8$ seconds. Numerical experiments show that it is feasible to use up to 15% of the rated fan power for regulation service to the grid, without noticeably impacting the building's indoor environment and occupants' comfort.

Another contribution of this paper is a method to estimate the baseline fan power. Accurate estimation of the baseline power, which is what would be observed in the absence of the regulation controller, is needed by the regulation controller as well as to assess the ancillary service delivered. In traditional demand response, estimating the baseline (sometimes called the counterfactual) is exceedingly challenging [17], [18]. We use signal bandwidth separation: the bandwidth of the baseline power is much lower than the bandwidth of the regulation signal. Since a commercial building has large thermal capacitance, its baseline fan power is a low frequency signal. In this paper, we constrain our frequency regulation signal to be a high frequency signal. As a result, the high frequency change in fan power, and thus air flow, will be absorbed by the building, and causing little change in the baseline power. Additionally, we are able to estimate this baseline power by using a low pass filter. This idea is in stark contrast to other demand response control strategies, which are mainly concerned with low frequency changes in demand, and have strong effect on the baseline power and its estimation accuracy.

A preliminary version of this paper was presented in [19].

It is important to notice the difference between what is proposed in this paper and demand response. While the latter is typically used for reducing power consumption of a load at emergencies, the proposed scheme involves both increase and decrease of power consumption. We therefore call the proposed methodology automated load tuning.

II. CONTROL ARCHITECTURE

A. Configuration of HVAC systems in commercial buildings

A typical HVAC system used in a modern commercial building, called a variable air volume (VAV) system, is shown in Fig. 2. Its main components consist of the air handing unit (AHU), supply fan, and VAV boxes. The AHU recirculates the return air from each zone and mixes it with fresh outside air. The ratio of the fresh outside air to the return air is controlled by dampers. The mixed air is drawn by the supply fan through the cooling coil in the AHU, which cools the air and reduces its humidity. In cold/dry climates it may also reheat and humidify the air. The air is then distributed to each zone through ducts. The VAV box at each zone has two actuators - a damper and a reheat coil. A controller at each zone, which we refer to as zonal controller, manipulates the mass flow rate of air going into the zone through the damper in the VAV box so that the temperature of the zone tracks a prespecified desired temperature, called zone set point. When the zone temperature is lower than the desired value, and flow rate cannot be reduced further due to ventilation requirements, the zonal controller reheats the supply air to maintain the zone temperature. As the zonal controllers change the damper positions in response to local disturbances (heat gains from solar radiation, occupants and so on), the differential pressure across the AHU fan changes, which is measured by a sensor. A fan controller changes the AHU supply fan speed, through a command to the variable frequency drive (VFD), so as to maintain the differential pressure to a predetermined setpoint. The VFD is a fast-responding and programmable power electronic device that changes the fan motor speed by varying motor input frequency. Since the air flow rate through the AHU is constantly changing to meet the demand from the zonal controllers, the system is called a VAV system. A complex interaction between a set of decentralized controllers and a top-level fan controller keeps the building at appropriate temperature while maintaining indoor air quality.

B. Proposed control architecture

The regulation signal sent by the grid operator is typically a sequence of pulses at 4 second intervals [20]. The magnitude of the pulse is the amount of deviation in their power consumption asked by the grid operator. Suppose the building is required to provide r(t) (in kW) amount of regulation service at time t. This signal is referred to as the (building-level) regulation reference. The job of a (building-level) regulation controller is to change the power consumption of the building so that the change tracks the regulation reference r(t).

We propose to achieve this through a feedforward controller. The power consumption of the supply fan is taken as the only source of flexibility. The controller changes the command to the fan so that the fan's power consumption is changed in such a way that the deviation in consumption – both positive and negative – tracks the regulation reference r(t). The architecture of the control system is shown in Fig. 3. The regulation reference r(t) is transformed to a *regulation command* u^r by the regulation controller. This command is then added to the baseline fan speed command u^b produced by the building's



Fig. 3. The proposed control architecture.

existing fan controller. Suppose $p^b(t)$ is the baseline power consumption of the fan due to the thermal load on the building, and $p^{b+r}(t)$ is the fan power consumption with the additional regulation command. Then the deviation in power consumed by the fan is $\Delta p(t) := p^{b+r}(t) - p^b(t)$. Clearly, changing the fan speed from its baseline value determined by the building's existing control system will change the air flow through the building. The goal is to design the regulation controller so that $\Delta p(t)$ tracks r(t) while causing little change in the building's indoor temperature.

Remark 1: In this paper, we assume that the power consumed by the boiler supplying hot water to the VAV boxes (for reheating) and the chiller/cooling tower providing chilled water to the cooling coil of the AHU are independent of the fan power. In many HVAC systems, the boiler consume natural gas instead of electricity. The second assumption may appear strong – the power consumed by the chiller and cooling tower may in fact change if the fan speed and, consequently, air flow rate changes. However, the dynamic interconnection between the AHU and the chiller can be thought of as a low pass filter due to the large mechanical inertia of the chiller/cooling tower equipment. Therefore, high frequency variations in the fan power will not change the power consumption of the chiller/cooling tower. Thus, the decoupling assumption - that fan power variations do not change chiller power consumption - holds as long as the variations are fast and of small magnitude. In addition, in some HVAC systems chilled water is supplied from a water storage tank. For such systems, the decoupling assumption holds naturally.

III. DATA-DRIVEN CONTROL ORIENTED MODELING

The dynamics of the complete closed loop system of a building relating zone temperatures to fan speed command are quite complex due to the interconnection of the zonelevel controlled dynamics, dynamics of pressure distribution in the ducts, and building-level fan controller. For the purpose of control design, we derive simplified models of some of these components using data collected from Pugh Hall on the University of Florida campus, a typical example of a modern commercial building.

A. Fan power consumption model

The power consumption of a fan is proportional to the cube of its speed [21]

$$p(t) = c_1(v(t))^3,$$
 (1)

where c_1 is a constant, and v is the normalized fan speed in percentage. For example, 100 indicates that the fan is running at full speed, and 50 means it is running at half speed. Additionally, the supply air flow rate is given by

$$m(t) = c_2 v(t), \tag{2}$$

where c_2 is a constant.

In practice, the fan speed is controlled by the VFD which also accelerates or decelerates the fan motor slowly in the interest of equipment life. Because of this ramping feature of VFD, we assume the transfer function from the control command to the fan speed is of first-order

$$\tau_1 \frac{dv(t)}{dt} + v(t) = u(t), \tag{3}$$

where τ_1 is the time-constant, and u(t) is the fan speed command sent by the fan controller. The fan speed controller is typically a proportional-integral (PI) controller. Note that both v and u are expressed in percentage.

In the AHU, the fan speed is controlled by the fan controller so that the actual fan speed v(t) tracks a desired fan speed $v^d(t)$. In practice the desired fan speed, $v^d(t)$, is communicated to the fan indirectly through a change in the duct pressure caused by the actions of the zonal controllers. For instance, when the desired mass flow rate $m^d(t)$ from the building increases, some dampers in the VAV boxes open wider, the duct pressure drops, and the supply fan needs to accelerate to maintain the duct pressure at a prespecified setpoint. Taking into account the air transportation in the duct, here we assume the transfer function from the desired mass flow rate $m^d(t)$ to the desired fan speed $v^d(t)$ is of first order

$$\tau_2 \frac{dv^d(t)}{dt} + v^d(t) = m^d(t)/c_2,$$
(4)

where τ_2 is the time-constant, and the devision of $m^d(t)$ by c_2 is due to Eq. (2). In this paper, we make the simplifying assumption that the fan controller senses the desired fan speed $v^d(t)$ directly and changes the fan speed to track this desired value. This allows us to sidestep the very challenging problem of modeling the duct pressure dynamics. Yet, the assumption is justified since that is what the fan control loop does, albeit indirectly.

Fan power model from data: We now estimate the parameters c_1, c_2 and τ_1 in the models (1)-(3) from data collected from Pugh Hall. The data used is from one of the three AHUs in the building with a 35-KW rated fan motor which supplies air to 41 zones. Using least-squares method and a randomly chosen 24 hour long data set, the parameters are estimated to be $c_1 = 3.3 \times 10^{-5} \ kW$, $c_2 = 0.0964 \ kg/s$, and $\tau_1 = 0.1 \ s$. Estimation of the time constant τ_2 in Eq. (4) is challenging, since in the current HVAC system, the duct pressure respond to the change of desired air flow rate in a closed-loop manner due to the closed-loop fan speed controller. In this paper, we make a heuristic estimation that the time constant τ_2 is approximately 10 seconds. In Section VI, we will conduct numerical experiments to check the robustness of the designed controller based on this heuristically estimated value when there is a model plant mismatch.

B. Aggregate building thermal model

In what follows, a simplified thermal model of the building based on the aggregate *building temperature* T(t), which can be thought of as the average temperature of all zones, is discussed. This simple but non-linear thermal model relates the total mass flow rate to the building temperature.

Consider the following physics-based lumped thermal model of the building ([22], [23])

$$C\frac{dT(t)}{dt} = \frac{1}{R}(T_{oa} - T(t)) + c_p m(t)(T_{la} - T(t)) + Q, \quad (5)$$

where C, R are respectively the thermal capacitance and thermal resistance of the building, T_{oa} is the outside air temperature, c_p is the specific heat of air, m(t) is the supply air flow rate, and the leaving air temperature T_{la} is the temperature of the air immediately downstream of the AHU.

For the purpose of designing a transfer function for the regulation controller in Fig. 3, we linearize the aggregate thermal dynamics (5). We define \tilde{T} and \tilde{m} as the deviations of the building temperature and supply air flow rate from their steady-state values T^* and m^* :

$$T = T^* + T, \quad m = m^* + \tilde{m}.$$
 (6)

Substituting (6) into (5), we obtain the linearized model of the aggregate building thermal dynamics:

$$\frac{d\tilde{T}}{dt} = -\frac{1+c_p Rm^*}{CR}\tilde{T} + \frac{c_p (T_{la} - T^*)}{C}\tilde{m},\tag{7}$$

where we have assumed T_{oa} and Q to be constant for the time under consideration. We use this assumption and linearization only for design and test the design through simulations with time varying signals and high fidelity nonlinear model in Section VI.

We next aggregate the effect of all the zonal controllers into one controller that we call the *building temperature controller*, by imagining that it computes the desired total mass flow rate $m^d(t)$ based on the difference between the desired building temperature T^d and actual building temperature T(t), and then signals the fan controller to provide this mass flow rate. Since each of the zonal controllers in commercial buildings are usually PI controllers, we choose the building temperature controller to be a PI controller as well. The input of the PI controller is the temperature deviation from its desired value \tilde{T} and its output is the desired air flow rate m^d .

IV. REGULATION BY FAN COMMAND MANIPULATION

We claim that commercial buildings can provide regulation service to the grid without causing discomfort to occupants so long as the bandwidth of the regulation signal is suitably constrained. The considerations in determining this bandwidth are discussed here along with the control strategy implemented to extract regulation service.

The bandwidth of the regulation signal sent to buildings should be chosen with the following factors taken into account. First, high frequency content in resulting regulation command u^r (see Fig. 4) is desirable up to a certain upper limit. Since



Fig. 4. Schematic representation of the interconnection between zone air flow request and the fan speed control architecture integrated with regulation.

the thermal dynamics of a commercial building have lowpass characteristics due to its large thermal capacitance, high frequency changes in the air flow cause little change in its indoor temperature. The statement is also true for individual zones of the building. Additionally, the VFD and fan motor have large bandwidth so that high frequency changes in the signal u^r lead to noticeable change in the fan speed and, consequently, fan power. Both effects are desirable, since we want to affect the fan power consumption without affecting the building's temperature. However, an extremely high frequency content in $u^{r}(t)$ is not desirable as it might cause wear and tear of the fan motor. Likewise, u^r should not have very low frequency content. Otherwise, even if the magnitude of u^r is small, it may cause significant change in cooling provided over long periods of time, which in turn can produce a noticeable change in the temperature of the building. Furthermore, a large enough change in the temperature will cause the zonal controllers to try to change air flow rate to reverse the temperature change. In effect, the building's existing control system will try to reject the disturbance caused by u^r . Being a feedback loop, this disturbance rejection property is already present in the building control system. If the controllers in the building (fan controller as well as the zonal controllers) do not have high bandwidth, they would not reject high frequency disturbance. In short, the frequency content of the disturbance $u^{r}(t)$ should lie in a particular band [f_{low} , f_{high}], where the gain of the closed loop transfer function from u^r to fan speed v is sufficiently large while that of the transfer function from u^r to temperature T is sufficiently small.

The parameters f_{low} , f_{high} are design variables to compute a suitable regulation signal for the building. We assume that these variables describing the bandwidth along with the total capacity of regulation that the building can provide is communicated to the grid operator or aggregator during deployment and used in constructing an appropriate regulation signal for the building. In the simulation studies described in Section VI, this is achieved in two stages. First, we pass the ACE data r(t) through a bandpass filter with a passband $[f_{\text{low}}, f_{\text{high}}]$.



Fig. 5. Magnitude vs. frequency of the closed loop transfer functions from disturbance to fan speed H_{v^Tv} , and from disturbance to temperature H_{v^TT} .

Second, we design the PI gains of the fan controller and zonal controllers so that the closed loop gain criteria described above are met. This in practice may be an iterative design process.

Suppose the regulation signal to be tracked by the building is $r^{\text{filt}}(t)$. As described in Section II-B, we need the fan power deviation $\Delta p(t) = p^{b+r}(t) - p^b(t)$ to track $r^{\text{filt}}(t)$. Let v^b be the baseline speed which would be observed in the absence of the regulation controller, we need to compute the additional fan speed v^r so that $c_1(v^b + v^r)^3 - c_1(v^b)^3 = r^{\text{filt}}(t)$ (see Eq. (1)). We compute the additional fan speed by using a first order Taylor series approximation: $r^{\text{filt}}(t) = 3c_1(v^b(t))^2v^r(t)$. Specifically, converter block in Fig. 4 is a static function that computes the command v^r

$$v^{r} = \frac{r^{\text{nn}}(t)}{3c_{1}(v^{b}(t))^{2}}.$$
(8)

The resulting command v^r is then passed through a prefilter to produce the command u^r . The fan speed command that is sent to the VFD is $u^b + u^r$. The prefilter is needed so that the gain of the transfer function from v^r to v in the band $[f_{\text{low}}, f_{\text{high}}]$ is close to 1, see Fig. 5. In this figure, as well as in simulation studies, we take $[f_{\text{low}}, f_{\text{high}}]$ to be [1/600, 1/8]. The choice of these values will be explained later. The prefilter is designed by computing an approximate inverse of the transfer function from u^r to v, which is calculated based on the diagram in Fig. 4 by replacing the dynamics of all individual zones and zonal controllers with the linearized aggregate building thermal model (7) and building temperature controller described in Section III-B. The magnitude responses of two crucial transfer functions are shown in Fig. 5. We see from the figure that within the prespecified band, the transfer function from disturbance v^r (regulation command) to fan speed v has a relatively high gain while to the temperature T has an extremely low gain.

A. Estimating the baseline on-line

To implement the control scheme described above, an online estimate of the baseline fan $v^b(t)$ speed is needed. The accuracy of the estimate has a strong effect on the tracking performance of the regulation signal. We use timescale separation to facilitate estimating the baseline fan speed. Since a building has low pass characteristics, the baseline fan speed is a low frequency signal. As long as there is enough separation between the frequencies of the regulation reference and the bandwidth of the building's closed loop HVAC control system, the fan speed commanded by the building's controller will not react to the changes to the fan speed commanded by the regulation command, and vice versa. In that case the low frequency content of the fan speed is what the building would have done without the regulation controller, i.e., the baseline fan speed. We therefore estimate the baseline fan speed by filtering the measured fan speed through a low pass filter with cut-off frequency lower than f_{low} . Fig. 4 shows the overall control architecture. The estimated baseline fan speed is denoted by \hat{v}^b , and the estimated baseline fan power is therefore given by $\hat{p}^b = c_1(\hat{v}^b)^3$, due to Eq. (1). The regulation controller uses $\hat{v}^b(t)$ instead of $v^b(t)$ in Eq. (8).

V. HIGH FIDELITY NON-LINEAR MODEL FOR SIMULATION STUDIES

Although a simplified thermal model is used for control design presented in Section III, we use a complex physicsbased model in the simulation studies aimed at testing the controller's performance. This model is briefly described next; see [23] for details.

To cope with the difficulty of modeling duct pressure dynamics that couple zone level dynamics to the fan dynamics, we make the following simplification. We assume that each zonal controller asks for a certain amount of air flow rate, by generating a *desired* air flow rate command $m_i^d(t)$ in response to the measured temperature deviation from the set point: $T_i^d(t) - T_i(t)$. The total desired supply air flow rate, $m^d(t)$, is the sum of the desired air flow rate into each zone $m_i^d(t)$:

$$m^{d}(t) = \sum_{i=1}^{n} m_{i}^{d}(t).$$
 (9)

The signal $m^d(t)$ is the input to compute the desired fan speed $v^d(t)$ (cf. Eq. (4)). The *actual* total mass flow rate is $m(t) = c_2v(t)$, where v(t) is the actual fan speed. It is divided among the zones in the same proportion as the desired air flow rates:

$$m_i(t) = \alpha_i m(t), \quad \alpha_i = m_i^d / (\sum_j m_j^d) . \tag{10}$$

The building's control system effectively performs this function, although signaling is performed through physical interaction, instead of through the exchange of electronic signals.

The thermal dynamic model of a multi-zone building is constructed by interconnection of RC-network models of individual zones and the corresponding zonal controllers. We consider the following RC-network thermal model for each zone in the building:

$$C_{i}\frac{dT_{i}}{dt} = \frac{T_{oa} - T_{i}}{R_{i}} + \sum_{j \in \mathcal{N}_{i}} \frac{T_{(i,j)} - T_{i}}{R_{i,j}} + c_{p}m_{i}(T_{la} - T) + Q_{i}, \qquad (11)$$

$$C_{(i,j)}\frac{dT_{(i,j)}}{dt} = \frac{T_i - T_{(i,j)}}{R_{(i,j)}} + \frac{T_j - T_{(i,j)}}{R_{(i,j)}}.$$
 (12)

A widely used control scheme for zonal controllers in commercial buildings is the so-called "single maximum". There are three operating modes in this control scheme: cooling mode, heating mode, and deadband mode. If all the zones are in the heating or deadband mode simultaneously, the supply fan will be maintained at a minimum speed to satisfy the ventilation requirement. However, this scenario is less common in practice. Additionally, our method changes the fan speed fast and with small magnitude. We therefore assume all the zones are in the *Cooling Mode*. In this mode, there is no reheating, and the supply air flow rate is varied to maintain the desired temperature in the zone. Typically a PI controller is used that takes temperature tracking error $T_i^d - T_i$ as input and desired air flow rate m_i^d as output.

The high fidelity model of a multi-zone building's thermal dynamics is constructed by coupling the dynamics of all the zones and zonal controllers, with m_i 's as controllable inputs, T_{oa} , Q_i , T_{la} as exogenous inputs, and T_i 's and m_i^d 's as outputs. The command m^d , computed using (9), is used to calculate the desired fan speed $v^d(t)$, which in turns serves as input to the fan controller, whose output is u^b . The total fan command $u^b + u^r$ is the input to the fan, with output fan speed v (which also determines the power consumption and mass flow rate through (1) and (2)). The mass flow rate through each zone, computed using (10), then serves as inputs to the building thermal dynamics. A schematic of the complete closed loop dynamics with the high fidelity model, along with all the components of the regulation controller, is shown in Fig. 4.

F-building test case: For our studies, we imagine a fictitious building with 4 stories and 44 zones, which we name *F-building*. Each story has 11 zones constructed by cutting away a section of Pugh Hall. Fig. 2 shows a layout of these 11 zones. The HVAC system of the F-building consists of a single AHU and zonal controllers for each of its zones. The F-building is meant to mimic the section of Pugh Hall serviced by one of the three AHUs that services 41 zones. The zones serviced by each of the AHUs in Pugh Hall are not contiguous, which necessitates such a fictitious construction. We identify the model of each of these 11 zones from data collected in Pugh Hall. Model identification consists of determining the R and C (resistance/capacitance) parameters in the model (11)-(12) for the zone. The least-squares approach described in [23] is used to fit the model parameters.

VI. REGULATION REFERENCE TRACKING BY F-BUILDING FAN

In this section, we describe simulation experiments which test the performance of the regulation controller described in Section IV for tracking regulation signal by varying the fan power. The F-building described in the previous section is used for the simulations reported here. The ACE signal r used for constructing the regulation reference r^{filt} is taken from a randomly chosen 4-hr long sample from PJM (Pennsylvania-New Jersey-Maryland Interconnection) [20]. It is then scaled so that its magnitude is less than or equal to 5 kW – a conservative estimate of the regulation capacity of F-building. A fifth-order Butterworth filter with passband [f_{low} , f_{high}] Hz is used as the bandpass filter while constructing r^{filt} . The design of the passband [f_{low} , f_{high}] will be discussed soon.

To unambiguously determine performance of the control scheme, we perform two simulations. First, a benchmark simulation is carried out with the regulation controller turned off so that $u^r(t) \equiv 0$. The fan speed is varied only by the



(a) Good tracking of a regulation signal with bandwidth $\left[\frac{1}{600}, \frac{1}{8}\right]$ Hz



(b) Poorer tracking of a regulation signal with bandwidth $\left[\frac{1}{1200}, \frac{1}{600}\right]$ Hz

Fig. 6. Numerical experiments of tracking regulation signals with different bandwidths.

building's closed loop control system to cope with the timevarying thermal loads. The observed fan power consumption - the true baseline - is denoted by $p^b(t)$. Second, another simulation is conducted with the regulation controller turned on and all the exogenous signals (heat gains of the building, outside temperature) identical to those in the benchmark simulation. The *actual* fan power deviation, $\Delta p(t)$, is defined as the difference between the fan power consumption observed in the second simulation, $p^{b+r}(t)$, and that in the first, $p^b(t)$. In addition, we define the *estimated* fan power deviation, $\Delta \hat{p}(t)$, as the difference between the measured fan power consumption in the second simulation $p^{b+r}(t)$, and the on-line estimate of the baseline power $\hat{p}^b(t)$ in the same simulation. Note that $\hat{p}^b(t) = c_1(\hat{v}_b(t))^3$ and $\hat{v}_b(t)$ is the output of the low-pass filter shown in Fig. 4.

The passband of the bandpass filter is designed based on tracking performance, which is measured by root mean squared tracking error. Simulations show that the regulation reference signal that can be successfully tracked by the proposed fan speed control mechanism is restricted in a certain bandwidth that is determined by the closed loop dynamics of the building. Fig. 6 (a) shows if the regulation reference $r^{\text{filt}}(t)$ has a bandwidth higher than 1/600 Hz (corresponding to period of 10 minutes), the actual fan power deviation $\Delta p(t)$ can track the regulation signal extremely well; the root mean squared tracking error is 0.12. Additionally, the estimated fan power deviation $\Delta \hat{p}(t)$ matches the actual fan power deviation $\Delta p(t)$ very well. However, if the regulation signal contains frequencies lower than 1/600 Hz, the zonal controllers compensate for the indoor temperature deviations in the zones by modifying air supply requirements, thus nullifying the speed deviation command of the regulation controller. This results in a poorer regulation tracking performance, as evidenced by Fig. 6 (b), in which the root mean squared tracking error is 0.49.



(b) Zone temperature deviations from set point

Fig. 7. Numerical experiment of tracking a regulation signal for a single building. The top plot (a) shows the baseline fan speed, fan speed with regulation, and online estimated baseline fan speed. The bottom plot (b) shows the temperature deviation T_i for each zone.

We estimate the upper band limit to be 1/8 Hz to avoid stress on the mechanical parts of the supply fan. In addition, since the ACE data from PJM is sampled every 4 seconds, the detectable frequency content in this data is limited to 1/8 Hz. Thus, the passband of the bandpass filter is chosen as [1/600, 1/8] Hz; cf. Fig. 5. Since we have neglected chiller dynamics, which has a time constant typically larger than 200 seconds [24], we expect that in practice the proposed control architecture will be able to successfully track regulation signal in the band $[1/\tau_0, 1/8]$, where $\tau_0 = 200$ (corresponding to period of about 3 minutes). At frequencies lower than that, unmodeled dynamics of chillers/cooling coils may affect performance.

When the bandwidth of the regulation signal is suitably constrained, we see that the actual fan power deviation tracks the regulation signal extremely well. Additionally, the baseline, estimated baseline and actual fan speed with regulation are depicted in Fig. 7 (a). We see that the estimated baseline speed matches the true baseline fan speed very all. The deviation of the actual fan speed with regulation control from baseline speed is small, which ensures that the pressure in the duct stays near designed values. Finally, Fig. 7 (b) depicts the temperature deviations of the individual zones from their set points. We observe that the temperature deviations are very small, which is unlikely to be noticed by the occupants.

In Section III, we made a heuristic estimation of the time constant τ_2 in Eq. (4) to be 10 seconds, and designed the regulation controller based on this heuristic estimation. We now test the designed controller (assuming $\tau_2 \equiv 10$ seconds) on systems whose actual time constant τ_2 varies between 5 seconds and 40 seconds to check the robustness of our designed controller when there is a model plant mismatch. We see from Fig. 8 that the designed controller is robust to uncertainty on the time constant estimation; the root mean



Fig. 8. Root mean squared errors of tracking with plant-model mismatch.

squared errors of tracking for system with different values of time constant are similar, and all of them are small.

A. Regulation potential of commercial buildings in the U.S.

The simulation results show that a single 35 kW fan can easily provide about 5 kW regulation capacity, which is about 15% of the total fan power. In Pugh Hall of University of Florida, there are two other AHUs, whose supply fan motors are 25 kW and 15 kW respectively. This means Pugh Hall by itself could provide about 11 kW regulation capacity to the grid. There are about 5 million commercial buildings in the U.S., with a combined floor space of approximately 72,000 million sq. ft, of which approximately one third of the floor space is served by HVAC systems that are equipped with VFDs [25]. Assuming fan power density per sq. ft. of all these buildings to be the same as that of Pugh Hall which has an area of 40,000 sq. ft., the total regulation reserves that are *potentially available* from all the VFD-equipped fans in commercial buildings in the U.S. are approximately 6.6 GW. Additionally, if we assume the regulation requirement of the United States is 1% of its peak load (such as in PJM), the regulation potential from fans is about 70% of the total regulation capacity currently needed [26].

VII. CONCLUSIONS AND FUTURE WORK

We showed that fans in commercial buildings alone could provide a large fraction of the current regulation requirements of the U.S. national grid without requiring additional investment. The proposed method is easy and inexpensive to deploy, which does not require additional sensors and actuators. A unique advantage of the proposed scheme is that the baseline power consumption of the building (that in the absence of the regulation controller) can be estimated on-line. This feature makes assessing the value of the ancillary service provided by the building extremely easy.

In this work we have neglected the effect of chillers and cooling coils on power consumption and indoor climate since their dynamics are slow. A more accurate characterization of the low frequency range at which regulation signal can be tracked with the proposed architecture requires incorporating these dynamics. In addition, chillers in a commercial HVAC system consume much more power than fans, and present a potentially larger opportunity. The problem of combined use of chillers and fans to provide ancillary service is the logical next step. Another avenue for further work is optimal dispatch of distributed energy resources by commercial buildings that have on-site distributed generation capability.

ACKNOWLEDGEMENT

The authors would like to express their appreciation to Johanna Mathieu for an inspiring discussion on estimating the baseline power of supply fan.

REFERENCES

- "Buildings energy data book." [Online]. Available: http:// buildingsdatabook.eren.doe.gov/default.aspx
- [2] B. Kirby, Frequency regulation basics and trends. United States. Department of Energy, 2005.
- [3] "Storage participation in ERCOT, prepared by the texas energy storage alliance," January 2010. [Online]. Available: http://www.ercot.com/
- [4] J. Smith, M. Milligan, E. DeMeo, and B. Parsons, "Utility wind integration and operating impact state of the art," *IEEE Transactions* on *Power Systems*, vol. 22, no. 3, pp. 900–908, aug. 2007.
- [5] U. Helman, "Resource and transmission planning to achieve a 33% rps in california-iso modeling tools and planning framework," in *FERC Technical Conference on Planning Models and Software*, 2010.
- [6] S. Meyn, M. Negrete-Pincetic, G. Wang, A. Kowli, and E. Shafieepoorfard, "The value of volatile resources in electricity markets," in *Proceedings of IEEE Conference on Decision and Control*, 2010.
- [7] K. Vu, R. Masiello, and R. Fioravanti, "Benefits of fast-response storage devices for system regulation in ISO markets," in *IEEE Power Energy Society General Meeting*, 2009, july 2009, pp. 1–8.
- [8] Y. V. Makarov, L. S., J. Ma, and T. B. Nguyen, "Assessing the value of regulation resources based on their time response characteristics," Pacific Northwest National Laboratory, Richland, WA, Tech. Rep. PNNL-17632, June 2008.
- [9] J. A. Short, D. G. Infield, and L. L. Freris, "Stabilization of grid frequency through dynamic demand control," *IEEE Transactions on Power Systems*, vol. 22, no. 3, pp. 1284–1293, 2007.
- [10] D. S. Callaway, "Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy," *Energy Conversion and Management*, pp. 1389–1400, 2009.
- [11] S. Kundu, N. Sinitsyn, S. Backhaus, and I. Hiskens, "Modeling and control of thermostatically controlled loads," in *the 17-th Power Systems Computation Conference*, 2011.
- [12] H. Hao, B. Sanandaji, K. Poolla, and T. Vincent, "A Generalized Battery Model of a Collection of Thermostatically Controlled Loads for Providing Ancillary Service," in *the 51th Annual Allerton Conference* on Communication, Control and Computing, 2013.
- [13] J. Braun, "Reducing energy costs and peak electrical demand through optimal control of building thermal storage," ASHRAE transactions, vol. 96, no. 2, pp. 876–888, 1990.
- [14] F. Oldewurtel, A. Ulbig, A. Parisio, G. Andersson, and M. Morari, "Reducing peak electricity demand in building climate control using real-time pricing and model predictive control," in 49th IEEE Conference on Decision and Control. IEEE, 2010, pp. 1927–1932.
- [15] D. Watson, S. Kiliccote, N. Motegi, and M. Piette, "Strategies for demand response in commercial buildings," in *Proceedings of the 2006* ACEEE Summer Study on Energy Efficiency in Buildings, August 2006.
- [16] S. Kiliccote, P. Price, M. A. Piette, G. Bell, S. Pierson, E. Koch, J. Carnam, H. Pedro, J. Hernandez, and A. Chiu, "Field testing of automated demand response for integration of renewable resources in californias ancillary services market for regulation products," 2012.
- [17] 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.
- [18] D. Claridge, "A perspective on methods for analysis of measured energy data from commercial buildings," *Journal of solar energy engineering*, vol. 120, no. 3, pp. 150–155, 1998.
- [19] H. Hao, A. Kowli, Y. Lin, P. Barooah, and S. Meyn, "Ancillary service for the grid via control of commercial building hvac systems," in *American Control Conference*, June 2013.
- [20] "PJM regulation data." [Online]. Available: http://www.pjm. com/markets-and-operations/ancillary-services/mkt-based-regulation/ fast-response-regulation-signal.aspx
- [21] ASHRAE, "The ASHRAE handbook HVAC systems and equipment (SI Edition)," 2008.
- [22] K. Deng, P. Barooah, and P. G. Mehta, "Mean-field control for energy efficient buildings," in *American Control Conference*, 2012.
- [23] Y. Lin and P. Barooah, "Issues in identification of control-oriented thermal models of zones in multi-zone buildings," in *IEEE Conference* on Decision and Control, December 2012.

- [24] Y. Yao, Z. Lian, and Z. Hou, "Thermal analysis of cooling coils based on a dynamic model," *Applied thermal engineering*, 2004.
- [25] "Commercial buildings energy consumption survey (CBECS): Overview of commercial buildings, 2003," Energy information administration, Department of Energy, U.S. Govt., Tech. Rep., December 2008.
- [26] "How much electric supply capacity is needed to keep U.S. electricity grids reliable?" [Online]. Available: http://www.eia.gov/todayinenergy/ detail.cfm?id=9671



He Hao received the B.S. degree in mechanical engineering and automation from Northeastern University, Shenyang, in 2006, the M.S. degree in mechanical engineering from Zhejiang University, Hangzhou, in 2008, and the Ph.D. degree in mechanical engineering from the University of Florida,

Gainesville, in 2012. He is currently a post-doctoral fellow in Department of Electrical Engineering and Computer Science, University of California, Berkeley. His research interests include smart grid, smart buildings, demand response, distributed control of large-scale multi-agent systems.



Yashen Lin received the B.S. degree in automation from University of Science and Technology Beijing, Beijing, in 2009, and the M.S. degree in mechanical engineering from University of Florida, Gainesville, in 2012. He is now pursuing a Ph.D. degree in mechanical engineering in University of Florida. His

research interests include smart buildings, building energy systems, and system identification.



Anupama S. Kowli (S'08-M'13) received the B.E. degree in electrical engineering from the University of Mumbai, India in 2006 and the M.S. and Ph.D. degrees in electrical and computer engineering from the University of Illinois at Urbana-Champaign in 2009 and 2013 respectively. Currently, she is an

assistant professor at Indian Institute of Technology Bombay. Her areas of interest include power systems planning and operations, electricity markets, and smart grids.



Prabir Barooah (S'03-M'02) received the B.Tech degree in mechanical engineering from the Indian Institute of Technology, Kanpur, in 1996, the M.S. degree in mechanical engineering from the University of Delaware, Newark, in 1999, and the Ph.D. degree in electrical and computer engineering from

the University of California, Santa Barbara, in 2007. He

is an Assistant Professor of Department of Mechanical and Aerospace Engineering, University of Florida, Gainesville, where he has been since October 2007. From 1999 to 2002, he worked at United Technologies Research Center as an Associate Research Engineer. His research interests include distributed control of energy systems for smart and green buildings, estimation in ad-hoc and mobile sensor networks, and cooperative control of autonomous agents.



Sean Meyn (M'87-SM'95-F'02) received the B.A. degree in mathematics from UCLA in 1982, and the Ph.D. degree in electrical engineering from McGill University in 1987 (with Prof. P. Caines). After 22 years as a professor at the University of Illinois, he is now Robert C. Pittman Eminent Scholar

Chair in the Department of Electrical and Computer Engineering at the University of Florida, and director of the new Center for Cognition & Control. His research interests include stochastic processes, optimization, complex networks, information theory, and power and energy systems.