Identification of Virtual Battery Models for Flexible Loads

Justin T. Hughes, *Student Member, IEEE,* Alejandro D. Domínguez-García, *Member, IEEE,* and Kameshwar Poolla, *Fellow, IEEE*

***Abstract*—The increasing prevalence of technologies such as advanced metering and controls and continuously variable power electronic devices are enabling a radical shift in the way frequency regulation is performed in the bulk power system. This is a welcome development in light of the increase of unpredictable and variable generation. The idea of active participation of loads in frequency markets is not new, but the rapidly changing landscape of the power grid requires new techniques for successful integration of new types of resources; this paper works towards that end. Previously, it has been shown that residential HVAC systems can be aggregated and used to provide frequency regulation by utilizing their thermal energy capacity and flexibility of energy consumption. The virtual battery model—a first-order linear dynamical model—was analytically shown to be an accurate and simple model to capture the flexibility of residential HVAC systems. This paper presents a technique for creating the same battery-type models for many other types of systems, which can be much more complex. Our technique is based on stress testing detailed software models of physical systems. A realistic case study involving the terminal building of a small airport is presented as evidence of the effectiveness of the proposed techniques.**

I. Introduction

IT has been hypothesized for decades that demand-side resources could supplement or even replace the regulation service provided by conventional generators [1]. The rapid changes that electric power systems are undergoing today have brought renewed interest in this idea. Intelligent communication and control devices and advanced metering are enabling technologies. Also, advances in power electronics allow more precise control over how much power loads consume, e.g., through cheaper and more reliable variable frequency drives. Lastly, the massive integration of renewable resources and their often unpredictable and variable generation patterns are creating an increasing need for frequency regulation [2]. Together, these transformative changes are creating a perfect environment for the participation of load side resources in ancillary service markets.

In this context, many recent papers have focused on the use of small residential thermostatically controlled loads (TCLs) with the ability to store thermal energy [3], [4], [5], [6]. Commercial buildings are another load with the ability to store thermal energy and with the flexibility to provide frequency regulation [7]. The large power consumption and thermal mass of these buildings, combined with the ability to continuously vary power consumption, could eliminate the need to aggregate them with other similar loads; this simplifies the communication and control with remote devices.

Frequency regulation is a necessary ancillary service because instantaneous imbalances between power supply and demand cause the frequency of a power system to deviate from its nominal value, e.g., 60 Hz in the USA. Many devices rely on a fixed frequency supply for proper operation, so this deviation must be minimized. System operators use this deviation alongside scheduled and actual interchange flows in the calculation of the area control error (ACE) [8]. This value is filtered and passed to a regulation controller [9], which generates the regulation signal which is automatically sent to generators every few seconds. The magnitude of the signal for each generator is determined by the fraction of its capability offer which cleared the market; for a more detailed description of the control center functions and market for the PJM interconnection, see [10]. Generators provide frequency regulation by adjusting their generation accordingly. Loads can also provide an equivalent service by adjusting their power consumption from what they otherwise would have consumed according to the regulation signal.

In order for a load to provide regulation, it must have a controller which can utilize its flexibility to follow a regulation signal. To this end, in this paper, we propose one such controller, the sole function of which is to provide maximum regulation capability while respecting constraints, e.g., occupant comfort and equipment ratings. We note that there may be additional objectives when designing the control system, e.g., minimization of total energy use [11]. A natural extension of our proposed controller would balance minimizing energy costs while maximizing income from the regulation market so as to minimize total costs.

Beyond designing the controller described above, in this paper we are interested in quantifying the flexibility of loads. Our tool for quantifying load flexibility is the virtual battery, which is useful because it is agnostic to the details of the underlying resource. It is worth noting that the use of virtual storage to refer to the ability of certain loads with the ability to store energy, not necessarily electrochemically form as in a regular battery, is gaining acceptance in the literature [12], [13], [14]. Sufficiently large loads will be able to offer capability into the ancillary services market directly as is done by generators and traditional storage devices. Others

J. T. Hughes and A. D. Domínguez-García are with the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, 61801. email: {hughes5, aledan}@illinois.edu

K. Poolla is with the University of California, Berkeley, Berkeley, CA, 94720. email: poolla@berkeley.edu

This work was supported in part by the U.S. Department of Energy under the Consortium for Electric Reliability Technology Solutions (CERTS).

may participate in a scheme where an aggregator makes a single bid on behalf of multiple smaller loads, then distributes the required regulation among them. In either case, it is easier to represent the resource with the battery model because it abstracts needless complexity for the application at hand. The authors of [3] demonstrated the power of this model to simply and succinctly describe the aggregate flexibility of a large number of possibly heterogeneous TCLs. In this paper, we show that the virtual battery model has the power to capture the flexibility of more complex loads, and we provide a method to identify its parameters.

Our proposed method for identifying the parameters of the virtual battery model requires a detailed model of the load and its control system. We use this model to perform software-based tests to determine equivalent battery parameters including charge/discharge rate limits, capacity, dissipation, and initial charge. These tests stress the system by issuing carefully selected commands to the controller. By noting which commands cause the controller to fail to converge at what time, our method can deduce how the real system will react to certain inputs; this is useful to determine how much frequency regulation can be offered in the market.

There is a small amount of literature in which techniques related to ours have been suggested. For example, in [15], charge rate limits and capacity parameters are identified for a collection of TCLs; our proposed method improves upon this technique in two ways. First, we identify the parameters in terms of a more accurate model that includes dissipation. Second, we do not rely on the ability to command a load to consume a maximum or minimum possible power. Although this may be simple for a collection of TCLs, it is not clear that this approach can easily be utilized in more complicated systems. The authors in [16] and [17] focus on using commercial buildings to provide regulation for certain frequency ranges of the regulation signal, but do not adopt the virtual battery model.

This paper expands and improves upon our previous work reported in [18]. The controller necessary to enable a load to provide frequency regulation has been revised to be numerically more robust, and it is now suitable for use in scenarios with heterogeneous parameters and constraints. The control inputs have been generalized to accommodate multiple types of controls. Together, these changes allow the proposed technique to be applied to a larger variety of loads including not only commercial building HVAC systems, but also other flexible loads. Our virtual battery identification procedure can now also identify the initial conditions (energy level) of the virtual battery model. Finally, we have performed a case study on the University of Illinois Willard Airport terminal building, which has a floor area of approximately $5900\,\mathrm{m}^2$. The end-use power consumption of this building peaks at over $600\,\mathrm{kW}$ in the late summer; the majority of this power is used for cooling. This study demonstrates the benefit and applicability of our work to realistic, complex systems and gives us some intuition as to the features that make for effective candidate buildings for providing frequency regulation.

The remainder of the paper is organized as follows. In Section II, we introduce the general model of the load and formulate our proposed controller using a first-order model as a motivating example. For clarity, we will continue to refer to the load as a commercial building with an associated HVAC system and controller; however, application to other systems is straightforward (e.g., by replacing the state vector with tank levels in a pumping application). In Section III, we will formulate the problem of identifying the parameters of the virtual battery model describing the ability of a load to provide frequency regulation. In Section IV, we introduce the proposed algorithm for identifying the parameters of the virtual battery model. Section V develops the airport model and presents the results of procedure used on this model. Concluding remarks are presented in Section VI.

II. Building Thermal Dynamics and Control

In this section, we introduce a generalized flexible commercial building model and a controller that enables it to provide frequency regulation. Throughout this section, we will use a first-order model to illustrate the ideas presented.

A. Flexible Load Model

We first introduce a model to describe the thermal dynamics of a building. Let T denote the vector of temperatures in different zones of the building, let s denote the vector of control inputs (e.g., HVAC fan speed, air flow control damper positions, on/off control for vestibule heaters, etc.), and let $w(t)$ denote a vector of exogenous variables (e.g. outside ambient temperature, solar thermal load, internal thermal load from occupants and equipment). Then, the dynamics of the system can be described by a nonlinear state space model of the form

$$\frac{d}{dt}T(t) = h_1(T(t), s(t), w(t)), \tag{1}$$

where $h_1(\cdot,\cdot,\cdot)$ describes the change in zone temperature as a function of the current state and control variables.

Additional variables relate the dynamics in (1) to the electric power consumed by the building HVAC system, which we denote $P(t)$; specifically, we can write

$$P(t) = h_2(T(t), s(t), w(t)). \tag{2}$$

Finally, we need to consider the constraints which arise from acceptable occupant comfort and those that arise from the ratings of the equipment:

$$h_3(T(t), s(t)) \leq 0. \tag{3}$$

Next, we illustrate the above concepts using a first-order model as an example.

Example 1 (First-Order System): Consider a single room building with heat transfer from outside ambient as well as interior thermal loads. Assume that the building is cooled by a variable speed air conditioner which recycles some fraction of the interior air. Let $T(t)$ be the temperature of the room and the control input $s(t)$ be equal to the mass flow rate of the conditioned air $\dot{m}$; then, the expression for (1) in this case is

$$\frac{d}{dt}T(t) = \frac{1}{m}\left(rT(t) + \dot{q} + c_p\dot{m}(t)(T_c - T(t))\right), \tag{4}$$

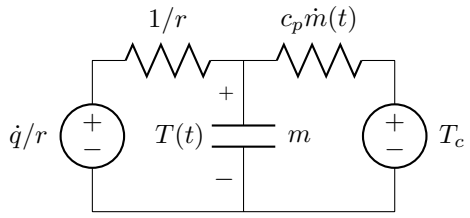

Fig. 1. Equivalent first-order model circuit diagram

where m is the thermal mass of the room, r is the thermal conductance, c_p is the specific heat capacity of air, and T_c is the temperature of the conditioned air. The external variable w equals $\dot{q}$, which is the thermal load that includes effects from the ambient temperature as well as objects inside the room. An equivalent circuit diagram of the thermal dynamics of the building is shown in Fig. 1.

The total power consumed, $P(t)$, is the sum of the fan power, which is assumed to be quadratic in mass flow rate, and the cooling power, which is the power required to maintain the cooling coil temperature as it cools the passing air; thus the expression for (2) in this case is

$$P(t) = \kappa_f \dot{m}(t)^2 + \frac{c_p}{\eta_c}\dot{m}(t)\left((1-d_r)T_{oa} + d_r T - T_c\right), \quad (5)$$

where κ_f is a function of the properties of the fan, η_c is the cooling system coefficient of performance, d_r is the fraction of return air that is recycled, and T_{oa} is the outside ambient temperature.

We enforce limits on interior temperature to ensure the comfort of occupants and limits on mass flow rate to keep the AC system functioning correctly; in this case, the constraints in (3) are as follows: $\underline{T} \leq T(t) \leq \overline{T}$, $\underline{\dot{m}} \leq \dot{m}(t) \leq \overline{\dot{m}}$. This completes the example.

B. Baseline Power

We define the regulation power at time t as the difference between the actual power consumed by the load, $P(t)$, and some baseline power, denoted by P^0, which is the total electric power consumed by the system were it not providing the regulation services. In practice, it may be challenging to calculate the baseline power in a way that is fair and auditable.

In this paper, we consider P^0 to be the value obtained from the steady state solution of (1), with the zone temperatures set to some nominal value, T^{m}. In subsequent developments, we will assume this solution satisfies (3); thus, by setting the left hand side of (1) to zero, we can implicitly solve for the baseline control input, s^0. If the problem is underdetermined, some secondary selection criterion would need to be utilized, e.g., energy efficiency. Using s^0, we can calculate the baseline power from (2), which results in the baseline power P^0.

Example 2 (First-Order System): Consider the same system as in Example 1; then, solving for the steady state of (4) gives us the baseline mass flow rate:

$$\dot{m}^0 = -\frac{rT^m + \dot{q}}{c_p(T_c - T^m)},$$

from which we can calculate the baseline power using (5); which results in

$$P^0 = \kappa_f(\dot{m}^0)^2 + \frac{c_p}{\eta_c}\dot{m}^0\left((1-d_r)\,T_{oa} + d_r T^m - T_c\right).$$

This completes the example.

C. Controller Design

The controller must enforce a constraint so the load consumes the commanded power output, $P^*(t)$, which is equal to the desired regulation plus the baseline power. The output of the controller is an optimal control, $s^*(t)$, which causes the HVAC system to consume the requested amount of power while also respecting the limits in (3). In the most general form, we can write

$$\begin{aligned} s^*(t) = \arg\min_{s(t)} \quad & h_4(T(t), s(t)) \\ \text{subject to} \quad & h_3(T(t), s(t)) \leq 0 \\ & |P^*(t) - P(t)| \leq \delta, \end{aligned} \quad (6)$$

where h_4 is an objective function that is used to weigh different possibilities if there are multiple solutions that satisfy the hard constraints and δ is some small value considered acceptable by the system operator.

The controller must be designed with the structure of the system and priorities of the owner in mind. For example, on a hot day the building HVAC system may be able to increase the amount of outside air being brought in. This would lower efficiency, causing extra power to be used without decreasing the temperature. This inefficient use of energy would result in a battery model with large capacity, but this would have to be balanced against wasting cool air on a hot day. An ideal objective function will have a unique solution that balances the capacity of the battery model, temperature variations, equipment cycling, and losses. The controller should also have defined behavior if it is unable to meet the commanded power.

Example 3 (First-Order System): For the first-order system in Examples 1 and 2, there is only one solution that satisfies the constraints in (6) if $\delta = 0$, so h_4 in (6) can be be chosen arbitrarily. In fact, we can find the solution analytically, so an optimization procedure is not required. To find this solution, note that at any given time we can find the mass flow rate that will cause a power consumption P^* by using the quadratic equation

$$P^*(t) = \kappa_f(\dot{m}^*(t))^2 + \frac{c_p}{\eta_c}\dot{m}^*(t)\left((1-d_r)\,T_{oa} + d_r T(t) - T_c\right)$$

and solving for $\dot{m}^*$. There will be only one meaningful solution assuming realistic parameters, which is given by

$$\begin{aligned} \dot{m}^*(t) = & \frac{c_p(-d_r T(t) + (d_r - 1)T_{oa} + T_c)}{2\eta_c\kappa_f} \\ & + \frac{\sqrt{c_p^2(-d_r T(t) + (d_r-1)T_{oa} + T_c)^2 + 4\eta_c^2\kappa_f P^*(t)}}{2\eta_c\kappa_f}. \end{aligned} \quad (7)$$

Plugging (7) into (4) gives us a an initial value problem which can only be solved numerically unless $d_r = 0$; this concludes the example.

III. Virtual Battery Parameter Estimation

We first define a procedure which incorporates software-based stress tests to determine which regulation signals the building/HVAC system is capable of following. Then, we introduce a reduced-order model—the virtual battery model—that we will use to compactly represent the flexibility of the building/HVAC system outfitted with the controller proposed in Section II-C. Using these, we formulate a criterion for the quality of the virtual battery model for describing the behavior of the full nonlinear system model. The problem is then to find the parameters that optimize this criterion.

A. Violation Time Function

We define a scalar input $u_i(t) = P_i^*(t) - P^0$ which is the desired deviation from the baseline power consumption profile. Assume we are free to choose $u_i(t)$, but have no knowledge of the structure or parameters of the underlying system in (1), (2), (3), and (6), and cannot make measurements beyond checking whether or not constraints have been violated. Then, for some input $u_i(t)$, define a function $f(u_i(t), \overline{\tau})$ such that if there is a constraint violation at or before $t = \overline{\tau}$, it takes the value of the time τ_i at which a constraint was violated, otherwise it takes the value ∞; in other words:

$$f(u_i(t), \overline{\tau}) = \begin{cases} \infty & \text{if } \exists \text{ solution to (6) } \forall\ t \leq \overline{\tau} \\ \tau_i & \text{otherwise,} \end{cases} \tag{8}$$

where $\tau_i = \min t$ such that there is no solution to (6).

B. Virtual Battery Model

The virtual battery been shown to accurately model the flexibility of certain buildings [3] [18]. Even in the case of a nonlinear, high-order building model, heat transfer is governed by Fourier's law. With the appropriate control, many buildings can act as a battery-type first-order model, the dynamics of which is given by:

$$\dot{x}(t) = -ax(t) - u_i(t), \tag{9}$$

where $x(t) \in \mathbb{R}$, $u_i(t) \in \mathbb{R}$, $a > 0$ is a constant, and $x(0) = x_0$. There are upper and lower bounds constraining $x(t)$ and $u_i(t)$, i.e.,

$$-C \leq x(t) \leq C, \quad -\underline{n} \leq u_i(t) \leq \overline{n}, \tag{10}$$

where $C > 0$, $\underline{n} > 0$, $\overline{n} > 0$ are constant. If a constraint is violated, the behavior is undefined. We group the parameters into a vector $\boldsymbol{\phi} = [a, C, \overline{n}, \underline{n}, x_0]^T$ to make the notation more compact.

For some input $u_i(t)$, define a function $b(u_i(t), \boldsymbol{\phi}, \overline{\tau})$ such that if a constraint in (10) is violated by (9) before time $\overline{\tau}$, it takes the value of the time τ_i at which a constraint was violated; otherwise it takes the value ∞. Thus, similar to (8), we have that

$$b(u_i(t), \boldsymbol{\phi}, \overline{\tau}) = \begin{cases} \infty & \text{if (9) – (10) hold } \forall\ t \leq \overline{\tau} \\ \tau_i & \text{otherwise,} \end{cases}$$

where $\tau_i = \min t$ such that (9) and (10) are not satisfied.

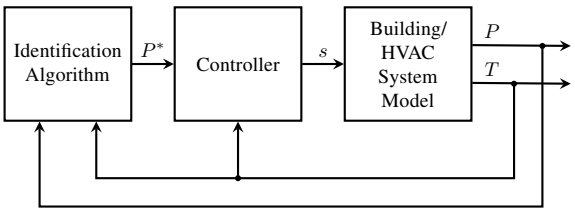

Fig. 2. System identification setup.

C. Problem Statement

We want to find the values of the virtual battery model parameters in (9) and (10) which will allow us to predict the behavior of the dynamic model in (1), (2), (3), and (6). The quality of the fit is inversely related to the difference between the violation times predicted by the nonlinear system and those predicted by the virtual battery model. If the fit is not exact, we wish to err on the side of caution by constraining the battery model to predict a violation time smaller than that incurred by the nonlinear model. This ensures that if an input does not cause a violation on the identified battery model, it will not not cause a violation in the nonlinear model. Mathematically, the problem can be formulated as finding a set of parameters $\boldsymbol{\phi}^*$ such that

$$\begin{aligned} \boldsymbol{\phi}^* = \underset{\boldsymbol{\phi}}{\text{argmin}} \max_{u_i(t)} \quad & |b(u_i(t), \boldsymbol{\phi}, \overline{\tau}) - f(u_i(t), \overline{\tau})| \\ \text{subject to} \quad & b(u_i(t), \boldsymbol{\phi}, \overline{\tau}) \leq f(u_i(t), \overline{\tau}). \end{aligned} \tag{11}$$

IV. Estimation Algorithms

In this section, we propose algorithms for identifying the parameters of the virtual battery model capturing the flexibility of the HVAC system of a commercial building as described by the dynamic model in (1), (2), (3), and (6). The basic structure of the proposed identification setup is shown in Fig. 2, where P^* is the commanded power, and s is a vector of control signals. Feedback includes the state vector T and the actual power P consumed by the building HVAC system.

A. Estimation of Rate Limits

The first step of the proposed procedure is to identify the rate limits $\overline{n}$ and $\underline{n}$. If the initial state is within its temperature bounds and we apply an input which causes a constraint to be immediately violated (i.e., $f(u_i(0), 0) = \infty$), we know it was due to the input constraints; this is because some finite time is required for an input to affect the value of the state.

To begin, we will assume w and $\boldsymbol{\phi}$ are constant. To be conservative, we must calculate $\underline{n}$ under the worst case scenario, which is when the state values are at their upper limit $\overline{T}$; likewise, $\overline{n}$ would need to use $\underline{T}$. The drawback of this conservative method is that the actual limits will be underestimated during normal operating conditions.

We know that $\overline{n} > 0$ ($\underline{n} > 0$), but we do not know an upper bound on these values. We therefore perform a one-sided binary search to find such an upper bound. Once we have an upper bound, we use it together with the greatest known lower bound in a binary search procedure to find $\overline{n}$ ($\underline{n}$) to arbitrary precision ϵ. The details of this procedure for estimating $\overline{n}$ and $\underline{n}$ are laid out in [18].

B. Estimation of Capacity and Dissipation Constant

If we respect the identified rate limits, we can guarantee that any constraint violation error is due to the capacity limit. In general, dissipation cannot be neglected when solving for the capacity limit so the two must be solved for simultaneously.

Because we are trying to fit the behavior of a linear model to that of a nonlinear one, we must look for a *sufficient* solution instead of an *exact* one. We say a solution is sufficient in the sense that verifying that an input does not cause any violations in the virtual battery model is sufficient to guarantee that the same input will not cause violations in the full nonlinear model. We are unable to mathematically prove that a solution is sufficient, i.e., the constraint in (11) will hold for all $u_i(t)$, because of the nonlinearity in (2); instead, we propose a heuristic. In Section V we will provide empirical evidence for the effectiveness of this heuristic procedure.

We start the procedure by picking a large value of $\overline{\tau}$ and generating $u_1(t), \ldots, u_n(t)$ such that violation times will be finite. Constant functions with different (log-distributed) magnitudes are a natural first choice, but others can be used. Then, we can construct a vector of violation times for the nonlinear model: $F = [f(u_1(t), \overline{\tau}), \ldots, f(u_n(t), \overline{\tau})]^T$. We also construct a vector of violation times for the linear model: $B(a, C, x_0) = [b(u_1(t), \boldsymbol{\phi}, \overline{\tau}), \ldots, b(u_n(t), \boldsymbol{\phi}, \overline{\tau})]^T$. The values of the components of the difference vector $B - F$ can vary by many orders of magnitude; in order to prevent the larger values from dominating the optimization, we define a new vector $G(a, C, x_0) = [\log|b(u_1(t), \boldsymbol{\phi}, \overline{\tau}) - f(u_1(t), \overline{\tau})|, \ldots, \log|b(u_n(t), \boldsymbol{\phi}, \overline{\tau}) - f(u_n(t), \overline{\tau})|]^T$, which is the natural logarithm of each component of $B - F$.

We wish to find the values of a, C, and x_0 so the two models have similar violation times. If the difference in times cannot be reduced to zero, we require that the violation time predicted by the virtual battery model be less than that of the nonlinear model. We can cast this problem as a constrained nonlinear least squares problem of the form

$$\begin{aligned} \underset{a,C,x_0}{\text{minimize}} \quad & \|G(a, C, x_0)\|_2 \\ \text{subject to} \quad & B(a, C, x_0) \leq F. \end{aligned} \tag{12}$$

If $u_i(t) = k$ is constant and $x(0) = x_0$, we can write an analytic expression for b. The solution to (9) is given by

$$x(t) = \left(x_0 + \frac{k}{a}\right) e^{-at} - \frac{k}{a}.$$

Then, by setting $x(\tau_i) = -C$ and solving for τ_i, we obtain

$$\tau_i = b(k, \boldsymbol{\phi}, \overline{\tau}) = \frac{1}{a} \log\left(\frac{-ax_0 - k}{aC - k}\right),$$

if $k > aC$. However, in general we cannot analytically solve for the first τ_i at which the equality given by

$$x(\tau_i) = x_0 + \int_0^{\tau_i} e^{-a(\tau_i - \tau)} u_i(\tau)\, d\tau = \pm C$$

holds. If this is the case, we can use a numerical method to evaluate the function b, just like we must use for f. Given x_0 and $\dot{x}(t) = -ax(t) - u_i(t)$, we can approximate the system trajectory, $x(t)$, $t > 0$, until $|x(\tau_i)| = C$, and record τ_i. This change makes the computation of B to take orders of magnitude longer, but we found that solving (12) is still computationally tractable. Because F needs to be calculated only once, we expect the procedure to scale reasonably well to large systems.

V. Case Studies

We begin this section by formulating a specific building/HVAC system model and a controller of the very general form presented in Section II. We will then test the performance of the identification procedure described in Section IV on the aforementioned system.

A. Commercial Building/HVAC System Model

Our building/HVAC system dynamic model and controller formulation is specifically tailored to describe the thermal dynamics of the University of Illinois Willard Airport. The formulation is a generalization of the commercial building model used in [18], which was found to be insufficient. Key changes include allowing for multiple air handling units, generalizing the fan power consumption formula, generalizing the control vector, and accommodating unconditioned zones.

Let T denote the vector of building zone temperatures; M denote the (diagonal) matrix of zone thermal inertias; T_{oa} denote the outside ambient temperature; c_p denote the specific heat capacity of air (assumed constant); $\dot{m}$ denote the vector of conditioned air mass flow rates to each zone; T_c denote the air temperature at the output of the cooling coils; $\dot{Q}_{\text{people}}$, $\dot{Q}_{\text{lighting}}$, $\dot{Q}_{\text{equipment}}$, $\dot{Q}_{\text{solar}}$ denote vectors of thermal loads due to occupants, electric lighting, miscellaneous machinery and office equipment, and solar radiation, respectively. Then, the dynamics of the building/HVAC system are described by

$$\begin{aligned} M\frac{d}{dt}T(t) =& R_1 T(t) + R_2(T_{oa} - T(t)) + \dot{Q}_{\text{people}} \\ &+ \dot{Q}_{\text{lighting}} + \dot{Q}_{\text{equipment}} + \dot{Q}_{\text{solar}} \\ &+ c_p \dot{m}.^*(T_c \mathbb{1} - T(t)), \end{aligned} \tag{13}$$

where $.^*$ indicates element-wise multiplication; R_1 is a sparse, symmetric matrix with off-diagonal elements (i, j) that are the thermal conductance between zones i and j, and diagonal elements (i, i) equal to the negative sum of the off-diagonal elements in row i; R_2 is a diagonal matrix of external conductance values.

By defining $R = R_1 - R_2$ and $\dot{Q} = R_2 T_{oa} + \dot{Q}_{\text{people}} + \dot{Q}_{\text{lighting}} + \dot{Q}_{\text{equipment}} + \dot{Q}_{\text{solar}}$, we arrive at a more compact version of (13):

$$M\frac{d}{dt}T(t) = RT(t) + \dot{Q} + c_p \dot{m}(t).^*(T_c \mathbb{1} - T(t)). \tag{14}$$

The Willard Airport terminal building has five air handling units (AHUs). Let $\dot{M}$ be the vector of mass flow rates through each of the air handler units. We assume ideal ducts, so each element can be obtained by summing the elements of $\dot{m}$ corresponding to the appropriate air handler. Then, the expression for power consumed by the air handling fans is a second-order function of mass flow rate:

$$P_f(t) = \kappa_{f2} \dot{M}(t)^T \dot{M}(t) + \kappa_{f1} \mathbb{1}^T \dot{M}(t) + \kappa_{f0}.$$

If we assume the return air is a weighted average of the conditioned zones, the expression for the cooling coil power is given by

$$P_c(t) = \frac{c_p}{\eta_c} \mathbb{1}^T \dot{m}(t) \left(T_m - T_c\right), \tag{15}$$

where the temperature of the input to the AHU is

$$T_m(t) = (1 - d_r) T_{oa} + d_r T_r,$$

and the return air temperature is

$$T_r(t) = \frac{\dot{m}(t)^T T(t)}{\mathbb{1}^T \dot{m}(t)}.$$

Fast chiller dynamics are assumed in (15), and fixed supply air temperature is assumed in (14) and (15). While we note there is other literature that assumes fast chiller dynamics [8], some papers use a first order time delay [10] or utilize only the fan power [19] to provide regulation. Extensive chiller modeling is outside the scope of this paper, but future work could classify typical equipment into classes where 1) our assumption effectively holds, 2) there is a noticeable, but acceptable lag in the chiller response, or 3) the response is delayed so much that is is no longer useful for frequency regulation.

The model is composed of zones which may be conditioned or unconditioned. Let T_1 be the vector of conditioned zone temperatures, and $\dot{m}_1$ be the vector of conditioned zone mass flow rates. Similarly, let T_2 be the vector of unconditioned zone temperatures, and $\dot{m}_2 = 0 \cdot \mathbb{1}$ be the vector of unconditioned zone mass flow rates.

The constraints will then be $\underline{T} \leq T_1(t) \leq \overline{T}$ and $\underline{\dot{m}} \leq \dot{m}_1(t) \leq \overline{\dot{m}}$. Additional constraints, $\underline{\dot{M}} \leq \dot{M}(t) \leq \overline{\dot{M}}$, limit the total flow through each AHU.

The optimization problem in (6) is chosen to be

$$\begin{aligned} s^*(t) = \arg\min_{s(t)} \quad & \left\| \left(T_1(t + \Delta t) - T_1^{\mathrm{m}}\right) ./ \left(\overline{T} - \underline{T}\right) \right\|_\infty \\ \text{subject to} \quad & \underline{T}_1 \leq T_1(t + \Delta t) \leq \overline{T}_1 \\ & \underline{s}(t) \leq s(t) \leq \overline{s}(t) \\ & \left| P^*(t) - P_f(t) - P_c(t) \right| \leq \delta, \end{aligned} \tag{16}$$

where $./$ indicates element-wise division and $s = \dot{m}_1$. The objective function is chosen so that the controller will attempt to keep all conditioned zones near their midpoint temperature. If the regulation signal pushes the system to its limits, the controller will bring all zones to their temperature limit at the same time, at which point there will be a constraint violation.

The problem in (16) is solved using sequential quadratic programming using the baseline control input as an initial condition. As overly restrictive constraints can cause problems with the numerical solver, if no feasible solution to (16) can be found with $\delta = 0.1\,\mathrm{W}$, δ is increased until a solution is found. Then, δ is decreased in gradual steps, using the previous solution as an initial condition for solving each new optimization problem. Then, if the optimization finds a solution with $\delta = 0.1\,\mathrm{W}$, the simulation continues. If there is truly no feasible solution, the controller considers this a violation and returns the solution that converged with the smallest δ. This alternative solution minimizes the error without violating state or input constraints. This graceful degradation is not required for the parameter identification algorithm proposed in Section IV, but it would be desirable when implementing the controller in a real system.

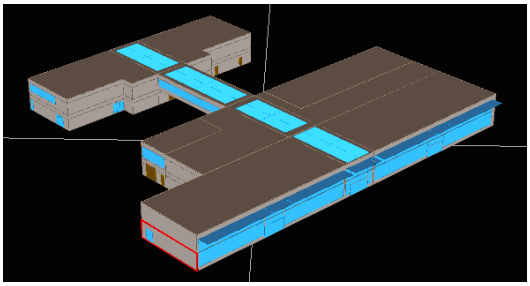

Fig. 3. Three-dimensional model of Willard Airport terminal provided by Shawn Maurer, Andrew Robinson, and Todd Rusk, all affiliated with the Illinois Smart Energy Design Assistance Center (SEDAC) at the University of Illinois [20].

To find the baseline power, (14) can be broken into

$$\begin{aligned} \begin{bmatrix} M_c & 0 \\ 0 & M_u \end{bmatrix} \begin{bmatrix} \dot{T}_1(t) \\ \dot{T}_2(t) \end{bmatrix} = & \begin{bmatrix} R_{11} & R_{12} \\ R_{21} & R_{22} \end{bmatrix} \begin{bmatrix} T_1(t) \\ T_2(t) \end{bmatrix} + \begin{bmatrix} \dot{Q}_1(t) \\ \dot{Q}_2(t) \end{bmatrix} \\ & + c_p \begin{bmatrix} \dot{m}_1(t) \\ \dot{m}_2(t) \end{bmatrix} .* \begin{bmatrix} T_c \mathbb{1} - T_1^m(t) \\ T_c \mathbb{1} - T_2(t) \end{bmatrix}. \end{aligned} \tag{17}$$

Then, we set $\dot{T}_1(t) = \dot{T}_2(t) = 0$, $T_1(t) = T_1^m$, and $\dot{m}_2 = 0$ in (17), giving us

$$\begin{bmatrix} R_{11} & R_{12} \\ R_{21} & R_{22} \end{bmatrix} \begin{bmatrix} T_1^m \\ T_2^0 \end{bmatrix} + \begin{bmatrix} \dot{Q}_1 \\ \dot{Q}_2 \end{bmatrix} = \begin{bmatrix} -c_p \dot{m}_1^0 .* (T_c \mathbb{1} - T_1^m) \\ 0 \end{bmatrix}. \tag{18}$$

Equivalently, (18) can be written as two equations:

$$R_{11} T_1^m + R_{12} T_2^0 + \dot{Q}_1 = -c_p \dot{m}_1^0 .* (T_c \mathbb{1} - T_1^m) \tag{19}$$

and

$$R_{21} T_1^m + R_{22} T_2^0 + \dot{Q}_2 = 0. \tag{20}$$

From (20), it follows that $T_2^0 = -R_{22}^{-1}(R_{21} T_1^m + \dot{Q}_2)$. By plugging this into (19), we obtain

$$R_{11} T_1^m - R_{12} R_{22}^{-1} (R_{21} T_1^m + \dot{Q}_2) + \dot{Q}_1 = -c_p \dot{m}_1^0 .* (T_c \mathbb{1} - T_1^m),$$

which can easily be solved for baseline mass flow rate, yielding

$$\begin{aligned} \dot{m}_1^0 = & (R_{11} T_1^m - R_{12} R_{22}^{-1} (R_{21} T_1^m + \dot{Q}_2) + \dot{Q}_1) \\ & ./ (-c_p (T_c \mathbb{1} - T_1^m)). \end{aligned}$$

B. Numerical Values for Willard Airport Terminal Building

A three-dimensional model of the Willard Airport terminal building in Savoy, Illinois, was created by the Illinois Smart Energy Design Assistance Center (SEDAC) at the University of Illinois [20]. The model, shown in Fig. 3, was created in eQuest, a software program designed to evaluate building energy performance [21]. The model comprises 41 zones, 19 of which are conditioned. Since it is a very detailed model of a real building, it makes for a compelling case study. We will next look at the development of some key parameters of the study; however, values for every parameter cannot be presented due to space restrictions.

1) Thermal Conductance: To construct the interior thermal conductance matrix, R_1, and exterior thermal conductance matrix, R_2, the material and geometry of the surfaces between each of the 41 zones and all the other zones and the exterior were considered.

2) Thermal Mass: One of the most important parameters for our study is the thermal mass M. The dimensions of each zone are known, so the volume of air can be easily calculated. The volume multiplied by the density and specific heat capacity gives us the thermal mass. We assume the air is at a constant $25\,^{\circ}\mathrm{C}$ temperature and 1 atm pressure for density calculations.

It is well known that solids have a higher specific heat capacity than gases; thus the walls, floors, and furnishings in a room usually have a higher thermal mass than the air. Because eQuest contains thickness, density, and specific heat values for the building materials, we can calculate their thermal masses. For interior walls, half the thermal mass was assigned to each of the two zones it separates. External surfaces can develop a significant temperature gradient across the insulation, making the effective mass smaller. To account for this, we divided the calculated mass by a factor of 2 to reach an effective amount.

For each zone, the thermal mass of the air is calculated. Then, the thermal mass of walls, floors, and ceiling is estimated by multiplying the square footage by the thermal mass per square foot of a representative zone.

3) Thermal Loads: Solar radiation values were recorded at the time of peak cooling load. These values were multiplied by a scaling factor to account for the time of day and day of year using standard insolation formulas [22]. However, this simplification does not account for the exact geometry of the building or the effect of cloud cover. Occupancy, lighting, and equipment loads are each updated hourly based on their regular weekday or weekend schedule.

4) Ambient, Limit, Supply Temperature: Ambient temperature (T_{oa}) data from 2013 in nearby Springfield, Illinois, is used due to its availability. Conditioned zone temperature limits were selected as $21.2\,^{\circ}\mathrm{C}$ and $23.2\,^{\circ}\mathrm{C}$; these values are consistent with normal variations in temperature without the regulating controller. Historical data shows some variation in supply temperature T_c, but an average of $15\,^{\circ}\mathrm{C}$ was selected.

5) Other Scalar Parameters: Fan parameters and the chiller efficiency were estimated using least squares estimation against eQuest timeseries data. The authors believe that, for the purposes of this study, this calibration adequately compensates for other approximations.

C. Estimation Procedure Results

We next present the results obtained by using the battery model identification procedure described in Section IV on the Willard Airport terminal building model. Three summer days were selected for the study. Numerical results are summarized in Table I. We will first analyze the results for the time-invariant studies, followed by the time-varying studies.

1) Charge Rate Limit Estimation: We find $\overline{n}$ varies from $138\,\mathrm{kW}$ to $239\,\mathrm{kW}$, and $\underline{n}$ varies from $46.8\,\mathrm{kW}$ to $148\,\mathrm{kW}$. The sum of these values is always approximately $286\,\mathrm{kW}$. This sum is driven by the difference between the upper and lower mass flow rate limits of the air handler units. The way this difference is distributed between the two rate limits depends on the baseline power. For example, on a cooler day, the relatively small baseline power results in $\overline{n} > \underline{n}$. Because of the asymmetry, this system could be more effectively utilized in a market setting that treats up and down regulation as two distinct services.

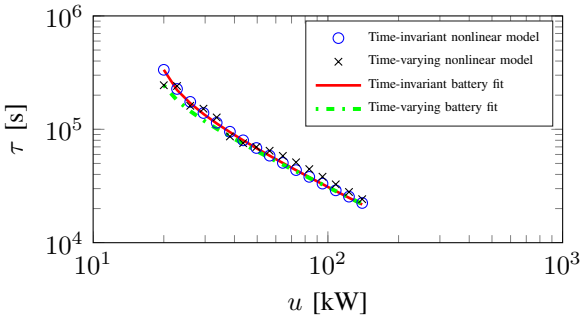

Fig. 4. Experimental violation times and best fit battery model predictions for simulations starting June 10th.

2) Capacity, Dissipation, and Initial Charge Estimation: The results in [18] show that for a time-invariant system results are independent of the type of input used. Step inputs are chosen for their simplicity.

Let $u_i(t) = k,\ t \geq 0$, with $k \in \mathbb{R}$. Sixteen values of k were chosen logarithmically-distributed between aC (which can also be found using a search procedure) and $\overline{n}$. Figs. 4 and 5 provide a plot of violation time versus input magnitude on June 10th and July 6th. If $a = 0$, we would expect a straight line with slope -1. The line curves upward for small inputs because there is more time for the effects of the dissipation to manifest themselves.

On June 10th, the parameters of the battery model obtained by our identification procedure are as follows: $a = 8.92 \times 10^{-6}\,\mathrm{s}^{-1}$, $C = 0.580\,\mathrm{MWh}$, and $x_0 = 0.228\,\mathrm{MWh}$. On July 6th, the identified parameters are $a = 7.60\,\mathrm{s}^{-1}$, $C = 0.312\,\mathrm{MWh}$, and $x_0 = 0.312\,\mathrm{MWh}$. We attribute the differences to changing environmental conditions (solar intensity, ambient temperature) and internal loads due to weekend versus weekday schedule. For both cases, even using our narrow range of acceptable temperatures, C is large compared to the charge rate limits. In fact, the virtual battery can supply maximum power for at least an hour without running out of charge, but a realistic regulation signal tends to alternate between charging and discharging with periods no longer than tens of minutes. Thus, we expect the charge rate limits to be the primary factor in determining the capability offer for the building.

D. Effect of Time-Varying Parameters

For the next set of studies, T_{oa}, $\dot{Q}_{\mathrm{people}}$, $\dot{Q}_{\mathrm{lighting}}$, $\dot{Q}_{\mathrm{equipment}}$, and $\dot{Q}_{\mathrm{solar}}$ are allowed to vary with time; in our previous work, these values were all assumed to be fixed. Fig. 6 shows how this translates into time-varying power values, and Fig. 7 shows how that creates diurnal patterns in temperature and mass flow rate. This generalization is challenging to our model, and analysis brings a number of interesting points.

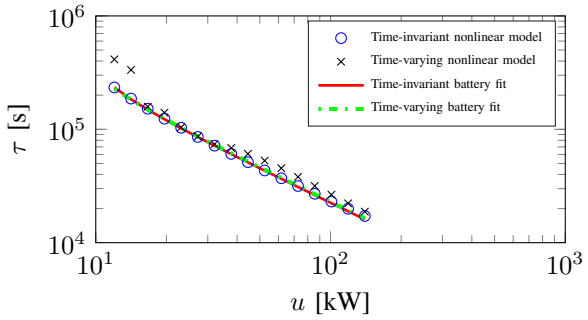

Fig. 5. Experimental violation times and best fit battery model predictions for simulations starting July 6th.

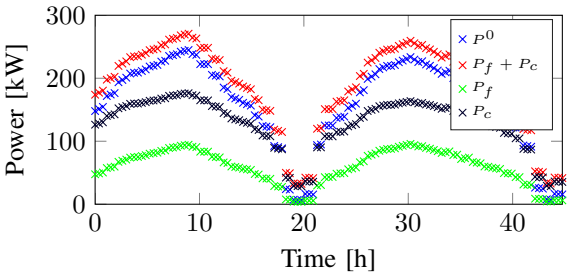

Fig. 6. Power values for the time-varying simulation starting June 10th.

First, we note for this study we assume knowledge of future values of the studied time-varying parameters. In practice, weather forecasts and historical thermal load data based on the regular airport schedule will need to be used. Future work could quantify the impact of uncertainty in predictions on output parameters. Further, as T_c and d_r may also vary, the procedure will need to be implemented in a receding-horizon manner, which will periodically update the identified battery parameters using new measurements to deal with the model-plant mismatch.

We note that the identified rate limits are identical to those in the time-invariant studies. This is because the rate limits are only identified at the first time step. As the baseline power

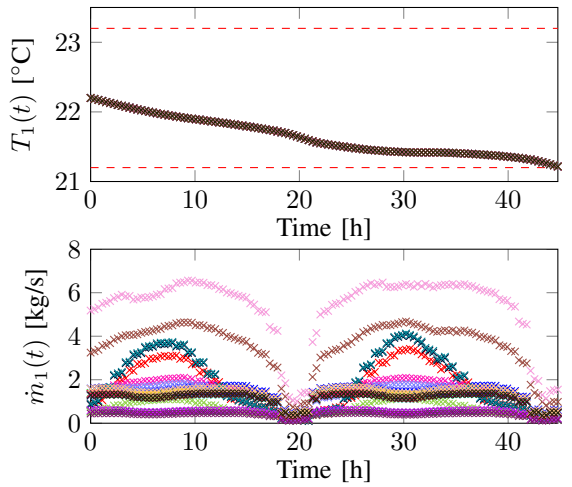

Fig. 7. Variables for the time-varying simulation starting June 10th.

TABLE I
IDENTIFIED PARAMETERS

Date	Time-varying	$\overline{n}$ (kW)	$\underline{n}$ (kW)	a (Ms^{-1})	C (MWh)	x_0 (MWh)
10-Jun	no	138	148	8.92	0.580	0.228
10-Jun	yes	138	148	11.7	0.424	0.424
06-Jul	no	239	46.8	7.60	0.312	0.312
06-Jul	yes	239	46.8	6.86	0.318	0.317
26-Aug	no	140	145	7.92	0.633	0.147
26-Aug	yes	140	145	7.06	0.708	0.0174

varies, the rate limits will also change. A simple fix is to generalize the battery model to allow time-varying rate limits and identify these rate limits at each time step of the simulation.

Next, we compare the parameters a, C, and x_0. It would be natural to also allow a and C to vary with time, but this would add significant complexity and make the model more difficult to analyze. We note that the time-varying battery fit curves in Figs. 4 and 5 do not fit the data from the time-varying nonlinear model as well as the time-invariant battery fit curves for the time-invariant nonlinear model data; thus, using the battery model would lead to more conservative predictions. The user would have to decide if the accuracy/complexity trade-off is appropriate for the task at hand.

We also notice that the time-varying and time-invariant fit curves in Figs. 4 and 5 are quite similar, but the identified parameters in Table I can differ significantly. This occurs because the battery parameters can be sensitive to small changes in the violation time data. The dissipation parameter a is most apparent over long periods of time. Also, the effects of C and x_0 are relatively indistinguishable to our method over short amounts of time. To distinguish them, the behavior over long periods of time is also critical. Unfortunately, it is over longer periods of time that the time-varying and time-invariant parameters tend to diverge. For some purposes, using the time-invariant battery model may be deemed sufficient to model a time-varying system. With sufficient data, studies can determine the functional relationship between the virtual battery parameters and exogenous variables such as time and ambient temperature, which can further improve the quality of the model.

E. Regulation Signal Input

In the final study, we illustrate the ability of a selected building to successfully follow a regulation signal that is within the capability characterized by the identified battery model without adversely impacting the indoor environment.

We utilize a normalized regulation signal $r(t)$, plotted in Fig. 8, which is chosen to be an example "regD" signal published by PJM. In this example, the identification procedure is performed each hour. In the following hour, the commanded deviation from the baseline power is $\overline{n}(t)r(t)$ if $r(t) \geq 0$, and $\underline{n}(t)r(t)$ if $r(t) < 0$. This scenario approximates participation in a real-time market with instantaneous hourly clearing of separate zero-cost up-regulation and down-regulation capability offers. If the identified capacity limit were not large compared to the charge rate limits, the offer would be based on the more restrictive limit.

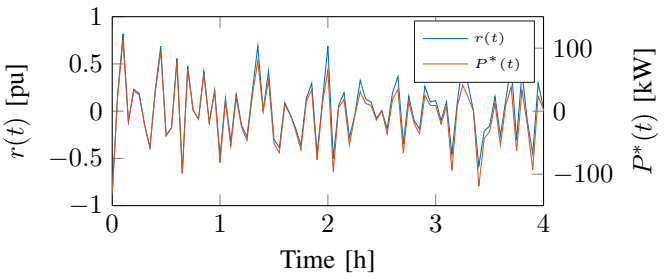

Fig. 8. Normalized regulation signal $r(t)$ and commanded deviation from baseline power $P^*(t)$. During this period, $\overline{n}$ and $\underline{n}$ are nearly equal, making the plots similar. During peaks and nadirs in power consumption, $P^*(t)$ is less balanced.

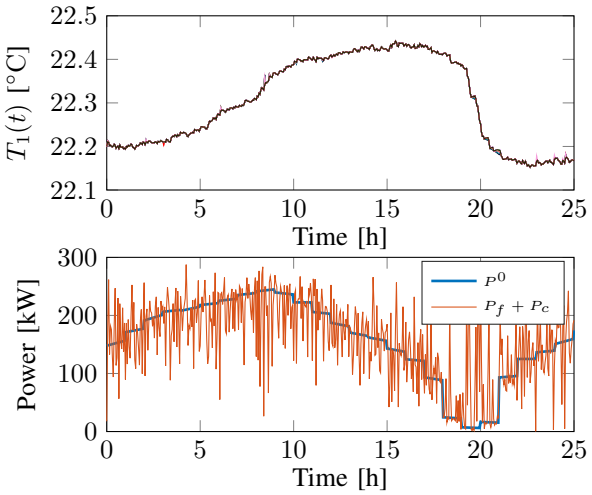

Fig. 9. Effect of tracking regulation signal on temperature and power. Temperature values remain well within the bounds of 21.2 °C and 23.2 °C. Simulation starting June 10th.

Fig. 9 shows how the controller proposed in Section II-C enables the building to track the regulation signal and the effect on indoor temperatures. The temperature variations are small compared to the established bounds of 21.2 °C and 23.2 °C. Power consumption varies greatly from the baseline, with more up-regulation at night and more down-regulation at midday. In this example, all constraints are satisfied at each time step. However, unless exogenous parameters are time-invariant, it is possible for situations to arise where this is not true if $r(t)$ takes on unfortunate values and external variable change quickly. Possible mitigation techniques include obtaining more accurate forecasts, identifying battery parameters more frequently, and making more conservative capability offers (e.g., 90% of calculated maximum) to account for uncertainty.

VI. Concluding Remarks

We have proposed a controller that allows for flexible loads to provide frequency regulation. We have introduced a method whereby the ability of the resulting closed-loop system to provide regulation can be accurately described by a simple, well understood battery model. Although the estimation method is an approximation, it was found to be effective on our University of Illinois Willard Airport test system.

Our case study revealed challenges that will be faced when applying the technique to real buildings. An ideal candidate building would have a large thermal mass with high, constant power consumption and the ability to consume much more or less power if required. Large thermal masses are common; however, arbitrary flexibility over power consumption is likely less abundant as excess capacity comes at a cost. Thus, for realistic buildings, rate limits will primarily be the limiting factor in determining regulation capability, and these parameters will vary with baseline power. Asymmetric regulation markets are extremely helpful in allowing full use of the virtual battery.

Acknowledgements

The authors would like to thank Shawn Maurer, Andrew Robinson, and Todd Rusk, all affiliated with the Illinois Smart Energy Design Assistance Center (SEDAC) at the University of Illinois, for providing us with the with the Willard Airport model data, fruitful discussions, and suggestions for improving the manuscript.

References

[1] F. Schweppe, R. Tabors, J. Kirtley, H. Outhred, F. Pickel, and A. Cox, "Homeostatic utility control," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-99, no. 3, pp. 1151–1163, May 1980.

[2] 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, May 2009.

[3] H. Hao, B. Sanandaji, K. Poolla, and T. Vincent, "Aggregate flexibility of thermostatically controlled loads," *IEEE Transactions on Power Systems*, vol. 30, no. 1, pp. 189–198, Jan 2015.

[4] B. Sanandaji, H. Hao, and K. Poolla, "Fast regulation service provision via aggregation of thermostatically controlled loads," in *Proc. of Hawaii International Conference on System Sciences*, Jan 2014, pp. 2388–2397.

[5] E. Vrettos and G. Andersson, "Combined load frequency control and active distribution network management with thermostatically controlled loads," in *Proc. of IEEE Conference on Smart Grid Communications*, Oct 2013, pp. 247–252.

[6] S. Meyn, P. Barooah, A. Busic, and J. Ehren, "Ancillary service to the grid from deferrable loads: The case for intelligent pool pumps in florida," in *Proc. of IEEE Conference on Decision and Control*, Dec 2013, pp. 6946–6953.

[7] H. Hao, A. Kowli, Y. Lin, P. Barooah, and S. Meyn, "Ancillary service for the grid via control of commercial building HVAC systems," in *Proc. of IEEE American Control Conference*, June 2013, pp. 467–472.

[8] A. Wood and B. Wollenberg, *Power Generation, Operation, and Control*. Hoboken, NJ: Wiley, 2012.

[9] A. Chakrabortty and M. Ilić, *Control and Optimization Methods for Electric Smart Grids*, ser. Power Electronics and Power Systems. New York: Springer, 2011.

[10] *Balancing Operations*, PJM, 2015, manual 12, revision 32.

[11] A. Kelman and F. Borrelli, "Bilinear model predictive control of a HVAC system using sequential quadratic programming," in *Proc. of IFAC World Congress*, Aug. 2011, pp. 9869–9874.

[12] G. Kats and A. Seal, "Buildings as batteries: The rise of "virtual storage"," *The Electricity Journal*, vol. 25, no. 10, pp. 59 – 70, 2012.

[13] H. Hao, A. Somani, J. Lian, and T. E. Carroll, "Generalized aggregation and coordination of residential loads in a smart community," in *Proc. of IEEE Conference on Smart Grid Communications*, Nov 2015.

[14] H. Hao and W. Chen, "Characterizing flexibility of an aggregation of deferrable loads," in *Proc. of IEEE Conference on Decision and Control*, Dec 2014, pp. 4059–4064.

[15] J. Mathieu, M. Kamgarpour, J. Lygeros, G. Andersson, and D. Callaway, "Arbitraging intraday wholesale energy market prices with aggregations of thermostatic loads," *IEEE Transactions on Power Systems*, vol. 30, no. 2, pp. 763–772, March 2015.

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

[17] 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," *IEEE Transactions on Smart Grid*, vol. 5, no. 4, pp. 2066–2074, July 2014.

[18] J. T. Hughes, A. D. Domínguez-García, and K. Poolla, "Virtual battery models for load flexibility from commercial buildings," in *Proc. of Hawaii International Conference on System Sciences*, Jan 2015, pp. 2627–2635.

[19] M. Maasoumy, B. Sanandaji, A. Sangiovanni-Vincentelli, and K. Poolla, "Model predictive control of regulation services from commercial buildings to the smart grid," in *Proc. of IEEE American Control Conference*, June 2014, pp. 2226–2233.

[20] "Smart energy design assistance center," [Online] http://smartenergy.illinois.edu/.

[21] U.S. Department of Energy Office of Energy Efficiency & Renewable Energy Building Energy Software Tools Directory, "eQuest," [Online] http://apps1.eere.energy.gov/buildings/tools_directory/software.cfm/ID=575.

[22] G. M. Masters, *Renewable and efficient electric power systems; 2nd ed.* Hoboken, NJ: Wiley, 2013.