

# Experimental Demonstration of Buildings Providing Frequency Regulation Services in the Swiss Market<sup>☆</sup>

Tomasz T. Gorecki<sup>a,\*</sup>, Luca Fabietti<sup>a</sup>, Faran A. Qureshi<sup>a</sup>, Colin N. Jones<sup>a</sup>

<sup>a</sup>*Automatic Control Laboratory, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland*

---

## Abstract

This paper investigates the potential of commercial buildings to offer ancillary services to the grid. We consider a tracking procurement inspired by secondary frequency control, where the resource needs to commit to a nominal consumption profile ahead of time and then receives a power consumption deviation around this profile. We propose a model-based predictive scheme which formally assesses the amount of power flexibility that can be offered by the resource and of attaining the best trade-off between comfort and financial gain during real-time operation by adjusting the nominal power consumption with a delay, in accordance to the regulations of the Swiss energy market. Extensive simulation results highlight how the possibility to perform adjustments on the scheduled energy profile is of paramount importance for both maintaining comfortable operational conditions and offering larger tracking flexibility. Finally, the effectiveness and robustness of the proposed methodology are demonstrated by means of full-day experiments in an occupied office space.

*Keywords:* Ancillary services; Demand response; Building control, MPC; frequency control; Power flexibility

---

## 1. Introduction

Power grids are undergoing massive changes to reach ambitious targets in terms of reduced carbon dioxide emissions, higher energy efficiency, economic competitiveness and increased security of supply. The increasing share of intermittent renewable energy sources connected to the grid challenges the current power grid stabilization paradigms. The increasing need for reserve power, which is now mostly provided by hydro-units and fast-ramping generation resources, has brought attention to the provision of regulation services by demand-side resources [1]. The potential of demand side resources has recently been identified by authorities in the USA and Europe. For example, the Pennsylvania-Maryland interconnection (PJM), has incentivized the participation of loads in ancillary services by adapting their participation rules [2]. Regulation services to the grid require accurate and fast control, therefore,

they usually take the form of direct control under the form of power production/consumption tracking [1], [3].

One challenge of demand side participation is that the loads have to maintain an appropriate quality of service while providing frequency regulation to the grid [1], which may be competing objectives. Nevertheless, numerous types of loads or pools of loads have been identified as suitable for providing regulation services, including thermally controllable loads [4], interruptible industrial and domestic loads [5], and plug-in hybrid electric vehicles [6]. [7] proposes a framework to study which resources are most suitable for each type of service. Building Heating, Ventilation and Air-Conditioning (HVAC) systems are more complex than most other systems considered since they are composed of numerous interacting subsystems and have many operational constraints. Numerous works have demonstrated the operational benefits of using predictive controllers for building systems [8], [9] for energy efficiency and cost savings. More recent work has concluded that savings might not always offset the cost of installation of the controllers [10]. Incorporating demand response strategies is however made possible by using such controllers and might further increase the savings [11].

The provision of regulation services follows the same logic and attracts a growing attention. Initial works have proposed methods to offer power consumption tracking using the HVAC system of a building [12], [13]. They assumed either a fixed regulation signal or estimated the capacity empirically. Several theoretical works have then proposed frameworks for computing the tracking capability

---

<sup>☆</sup>This work has received support from the Swiss National Science Foundation under the GEMS project (Green Energy Management of Structures, grant number 200021 137985), the NRP 70 Energy Turnaround Project (Integration of Intermittent Widespread Energy Sources in Distribution Networks: Storage and Demand Response, grant number 407040 15040/1), and from the European Research Council under the European Unions Seventh Framework Programme (FP/2007-2013), ERC Grant Agreement 307608.

\*Corresponding author

*Email addresses:* tomasz.gorecki@epfl.ch (Tomasz T. Gorecki), luca.fabietti@epfl.ch (Luca Fabietti), faran.qureshi@epfl.ch (Faran A. Qureshi), colin.jones@epfl.ch (Colin N. Jones)

ity of such systems. Several model-based methods using robust programming concepts have been proposed including [14], [15], [16], [17]. Other works have proposed analytical methods to analyze the aggregate flexibility of populations of loads in order to characterize it as an equivalent ‘virtual battery’ [4], [18], [19], [20]. A detailed description of these contributions is given in [21]. Aggregation of multiple buildings is considered in [22] and [23]. Experimental works focus more on the technical feasibility of simple strategies to implement power tracking in commercial HVAC systems. [2] proposes two methods to modulate the power consumption of a commercial HVAC system by either adjusting the fan duct pressure or the zone temperature setpoint (which indirectly influences the consumption of the fan) and tests those methods in simulations. In [24, 15], authors empirically estimate the power flexibility of a building. The power consumption of the main ventilation system is varied acting on the supply duct static pressure setpoint. Experiments were conducted showing how the resulting fluctuating air mass flow does not affect occupants’ comfort. In [25], power tracking is provided by adjusting fan power consumption. The building receives a filtered version of the tracking signal which is tracked with respect to the nominal baseline computed by the pre-existing building controller. Experiments were conducted on a 40,000 sq. ft. building over a 40 minute duration. [26] showcases the control of the variable speed compressor of a chiller through manipulation of the cooling water setpoint temperature. It is shown that the tracking quality requirements of PJM are met. [27] presents more extensive regulation experiments with a formal computation of the regulation capacity. Experiments are made in single-zone unoccupied test cells equipped with a standard cooling system.

Our previous work [28] demonstrated the applicability of the methods for calculating regulation participation described in [29] and [21] by means of laboratory-scale experiments using electric heaters. It differed from other published works by computing the optimal amount of frequency service offered rather than finding it through trial and error procedures. Moreover, it provided a methodology to compute the baseline ahead of time at the beginning of the experiments, therefore complying with the requirements of the Swiss electricity market.

*Notation:* Throughout the article,  $\mathbb{R}^l$  denotes the  $l$ -dimensional real space, upper case letters are used for matrices and lower case for vectors.  $a_i$  represents the value of vector  $a$  at time  $i$  whereas, bold letters are used to denote sequences over time, e.g.,  $\mathbf{p} = [p_0^T, p_1^T, \dots, p_{N-1}^T]^T$ .

### 1.1. Contribution of the paper

This paper demonstrates how to optimally use the thermodynamic storage of commercial buildings to offer frequency regulation services to the grid. The contribution with respect to our previous work [28] is twofold: First, we provide a method to determine the amount of regulation that can be provided taking into account the possi-

bility to adjust power consumption on the intraday market. Second, we show how to optimally trade regulation commitment and comfort in the building during real-time operations.

The control methodologies are validated in a series of experimental tests in a living laboratory. The experimental demonstration extends existing experimental works presented in the literature in the following ways:

- The method proposed is in full accordance with the Swiss regulation for secondary frequency control. In particular, the baseline consumption is determined a-priori as a result of the day-ahead market trades. Any modification of the baseline during the day respects current rules for the intraday market.
- Modeling of the influence of outside temperature and the sun was performed for these offices that are largely influenced by these factors. Weather forecasts were incorporated in the optimization to improve the previously proposed method.
- Experiments have been performed over extended periods of time (18 to 24h) in occupied offices. Experiments were successful despite large uncertainties in weather prediction and occupation and, therefore, demonstrate the robustness of the approach.

## 2. Problem Statement

### 2.1. The Swiss Energy market

The case study and the experiments undertaken in this work have been tailored to the Swiss regulatory framework of frequency regulation provision and energy dispatch [30]. Note that the guiding principles underlying the structure of the market are similar in all geographic regions but noticeable differences exist between markets. The main aspects of interest of the Swiss market are detailed in this section.

The energy procurement for Switzerland is done in a day-ahead market. Bids for energy trading are cleared at 11am on the day prior to delivery. Additionally, an intraday market allows participants to trade energy in 15 minute slots up to 60 minutes before delivery.

Ancillary services procurement are services traded between the network operator and power consumer or producer that aim at ensuring a safe operation of the power network. Frequency control is one of the categories of ancillary services. It is divided into three categories, namely primary, secondary and tertiary frequency control. The focus of this manuscript is secondary frequency control. The basic principle of secondary frequency control is to have providers change their power consumption in real time by tracking a signal called the Area Generation Control (AGC) signal. This service is rewarded by a payment proportional to the magnitude of the deviations the provider is willing to support. This maximum magnitude is agreed

Symbol	Description
$\gamma$	Capacity bid
$\bar{p}_{i DA}$	Day-ahead baseline purchase
$\bar{p}_{i i-\delta}$	Intraday transaction performed at time $i - \delta$ for baseline at time $i$
$\bar{p}_i$	Total Baseline Power Consumption
$\delta$	Minimum lag for intraday transactions
$y_i$	Outputs (Zone Temperatures)
$\epsilon_i$	Tracking error
$p_i$	Total Power Consumption
$u_i$	Control action
$d_i$	Weather disturbance at time $i$
$\tilde{\mathbf{d}}$	Weather forecast
$a_i$	Normalized AGC signal at time $i$
$\tilde{\mathbf{a}}$	AGC forecast

Table 1: Control architecture nomenclature.

upon in advance and embodied by a quantity called capacity bid, measured in MW. The main challenge when using loads for these kinds of services is the requirement to modify their power consumption around a fixed baseline on a fast timescale and for extended periods of time.

## 2.2. Control Structure

We consider a power consumer offering secondary frequency service. According to the Swiss market regulations, a tailored control architecture is proposed. A schematic of the controller structure is provided in Figure 1.

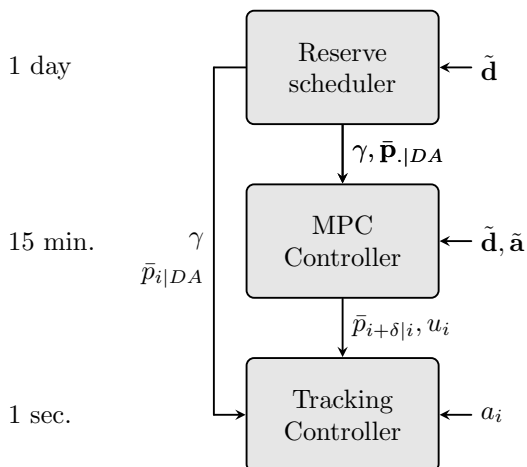


Figure 1: Architecture of the control system for tracking service procurement with participation in the intraday market. See Table 1 for the definitions of symbols.

Table 1 reports the nomenclature of the paper. The reserve scheduler decides on the capacity bid  $\gamma$  and the baseline purchased on the day-ahead market  $\bar{p}_{i|DA}$  for each time slot  $i$  of the following day. Subsequently it can readjust its power baseline up to one hour in advance on the intraday market by placing an order  $\bar{p}_{i|i-\delta}$  at time  $i - \delta$  for each time slot  $i$ . For simplicity, we assume that all intraday adjustments are made at time  $i - \delta$ , the closing time

for intraday transaction for time slot  $i$ . The final baseline is  $\bar{p}_i = \bar{p}_{i|DA} + \bar{p}_{i|i-\delta}$ . During the operation, a normalized AGC signal  $a_i$  is scaled by the bid  $\gamma$  and dispatched at a rate of one second for tracking to the building, which has to adapt its power consumption such that

$$|\epsilon_i| = |p_i - \bar{p}_i - \gamma a_i| \leq \alpha \gamma \quad (1)$$

where  $\epsilon_i$  is the tracking error. Note that the tracking error is allowed to scale in proportion to the capacity bid, in accordance with tracking regulations requirements as per [31]. It is readily seen that the problem has multiple timescales interacting: when deciding the capacity bid and day-ahead purchase, a prediction over a minimum of one day is required, while during operation, fast rate tracking is required. Similarly to [27] and [28], the control architecture proposed is therefore hierarchical with three interacting layers as depicted in Figure 1:

- The reserve scheduler commits the capacity bid for the next time period (one day) and buys energy on the day-ahead market. At the start of each day, it computes a capacity bid using the current estimate of the state of the system and up-to-date weather forecasts. As a result of this computation, the bid  $\gamma$  and day-ahead purchases  $\bar{\mathbf{p}}_{\cdot|DA}$  are committed. It is important to notice that the scheduler already takes into account the possibility to adjust the baseline later on on the intraday market, as will be detailed in section 4.
- A closed-loop Model Predictive Controller (MPC) operates at a fifteen minute timestep: its purpose is two-fold: compute adjustments to the baseline using the intraday market, and recompute optimal inputs for the heating systems based on updated forecasts and the current value of the AGC signal. As a result, an intraday trade is placed to adjust the baseline one hour ahead for a value  $\bar{p}_{i+\delta|i}$ . The inputs to the system for the upcoming time slot  $u_i$  are passed down to the fast controller.
- A fast controller modulates the power consumption of the HVAC at a fast rate to provide the tracking service. Based on the current received value of the AGC  $a_i$  and the committed baseline  $\bar{p}_i$ , the power input of the HVAC is controlled to meet the tracking requirement.

**Remark 1.** The fact that the baseline is decided one day in advance is a major difference with existing experimental works: in [25], the signal to track is filtered to keep only fast frequencies (spectral content below 10 minutes). Therefore, they propose to reconstruct the baseline from the final power consumption by filtering it with a low pass filter. Since this high-frequency deviation does not disturb the operation of other parts of the system, no rebound effect should be observed on the baseline. A similar idea is proposed in [32] where the baseline is determined




Figure 2: Floor map of the offices used as experimental testbed. The offices shaded in grey have been equipped with sensors and actuation and used for experiments. Red squares show the position of the heating units.

a-posteriori. Although that scheme solves a number of practical issues, it is currently not in full accordance with regulations, in contrast to our approach.

**Assumption 1.** The intraday markets are assumed to be liquid, meaning that it is always possible to sell or purchase energy according to the market clearing regulations.

In addition to experiments carried out during the day in occupied offices and for extended periods of time, the main difference in the formulation with our previous work [28] is the possibility to readjust the baseline in the intraday market. We will see that this feature is paramount to the robustness of the scheme. As was outlined in the introduction, demand-side resources are subject to sustained prediction errors that hinder the ability to predict accurately their power consumption on long time horizons. Therefore, the possibility to partly reschedule baseline is fundamental. We will see that the intraday market allows us to successfully meet the combined competing objectives of maintaining comfort and providing accurate frequency regulation. Clearly, this statement depends on the predictive power of our models, but the experiments demonstrate that for our system, which is affected by large uncertainties, the statement holds, which is a significant improvement over our preliminary work, where the level of uncertainty was smaller since experiments were conducted at nighttime.

### 3. System Modelling

#### 3.1. Testbed Description

The experiments were conducted in five office rooms accounting for a total area of 115 sq. meters ( $\sim 1200$  sq.

ft.). As depicted in Figure 2, each office is directly exposed to the outside and has varying exposure to the sun. Offices are labeled according to their exposure to the sun. Rooms NW, N, SW and S are individual offices. Room SE is a shared office occupied by six PhD students. During the experimental phase, all offices were occupied except for office SW. Each room is equipped with a commercial electric heater. The electric heaters were chosen for the following reasons: 1) Modelling electric heaters is simple 2) It is possible to measure their power consumption directly 3) They are highly responsive elements which permits the variation of their power consumption very quickly. This is a key element to offer frequency regulation services as fast continuous control is required. 4) Electric heating represents a significant share of the heating provision in Switzerland. Recent federal statistics indicate the presence of a quarter million electric-based heating units accounting for 4% of the total Swiss electricity consumption [33]. For measurements, wireless temperature sensors were used. Temperature sensors send measurements asynchronously to a server with a maximum frequency of six minutes. Weather data were collected from a nearby weather station every 30 seconds, including measurements for outside temperature, horizontal global solar radiation and weather observation. All the data are uploaded into a database on the local network.

The heaters are rated at 1900 Watts at 230 Volts, summing up to a total maximum power capacity of 7400 Watts. The heaters are normally equipped with a thermostat and a switch to adjust the level of heating between three distinct levels. In order to be able to modulate their power consumption continuously, the heaters were customized with additional hardware that allows pulse-width modulation (PWM) at 4 Hz. The control input to the system is the pulse-width modulation ratios denoted  $u$ . Therefore  $u \in [0, 1]$ . A pulse-width ratio directly results in an electric power consumption  $p = uP_{\max}$ . Finally, the heaters being resistive elements, the power consumption directly translates into a thermal power input to the room  $q$ , so that  $q = p = uP_{\max}$ .

Power modulation and communication at the heaters are managed by custom micro-computers (BeagleBone Black). The heaters receive setpoints from a server running on a regular desktop computer, and executing the MPC. The server collects temperature measurements from the web database, decides on the allocation of the AGC signal among heaters and communicates the power reference to the microcomputers, through Ethernet. More details regarding the experimental setup can be found in [28].

#### 3.2. Identification

In [28], models for the experimental testbed had been identified using standard black-box linear system identification. The model was describing the effect of the heater power input on the temperature of each room. It did not include the influence of the outside weather but featured a good predicting capability for nighttime experi-

ments. Since, we extended these first experiments to perform day-time experiment, the first step was to build a model that includes the influence of external perturbations from weather. The previously identified model was implicitly valid for specific average nightly outdoor conditions. Therefore the identification procedure was started from scratch again taking all factors into account.

For each zone, new experiments were conducted with open-loop input signals to excite the system. A mix of step tests and pseudo-random binary signals (PRBS) have been used in order to excite a wide range of frequencies in the system [34]. Of course, weather cannot be controlled but over the total duration of the identification experiments (about a week), a reasonably varying patterns of weather has been observed. Notice that weather is mostly slowly varying so we mainly hope to identify the static gain between the outside temperature, the solar gains and the indoor temperature. Each room was identified separately since the thermal coupling between them is weak.

For each room, the acquired data includes temperature measurements, power consumption of the heater, outside temperature, horizontal solar radiation, and cloud coverage. A model was identified for each zone with three inputs and one output, the room temperature. The inputs are power consumption of the heater, outside temperature and one input capturing the effect of the sun. The sun inputs for each zone are different. They are computed as the resulting total solar radiation on a surface facing one of the cardinal directions in  $\text{kW}/\text{m}^2$  (for example a west facing surface for room NW). Global horizontal radiation rate and weather description are collected, from which radiation rates on different surfaces are reconstructed. The reader is referred to [35], [36] for detailed information on solar radiation reconstruction.

Each zone was modeled with a second order Auto Regressive model with exogenous inputs (ARX) [34]: the inputs of the model are  $q$ , the thermal power input and  $d = [T_o, q_{sun}]$ , the vector of disturbance affecting the system, namely the outside temperature and the solar radiation input and the output the room temperature  $y$ .

$$A(r)y(t) = B_1(r)q + B_2(r)d(t-1) + e(t) \quad (2)$$

where  $A(r)$ ,  $B_1(r)$  and  $B_2(r)$  are polynomials of the delay operator  $r$  and  $e(t)$  a white noise disturbance. The coefficients of  $A$ ,  $B_1$  and  $B_2$  are found by solving a least square estimation problem using the input and output data recorded during the experiments [34].

Figure 3 reports the prediction of the model against the measured output for room NW for one of the identification experiment. The normalized root mean square prediction error for the models ranges between 70 and 90 % depending on the room. In particular, room N has a good model since it is not directly exposed to the sun while room SW has the model with lowest fit because it is difficult to capture the impact of the sun accurately since the room has a facade facing South and one facing West.

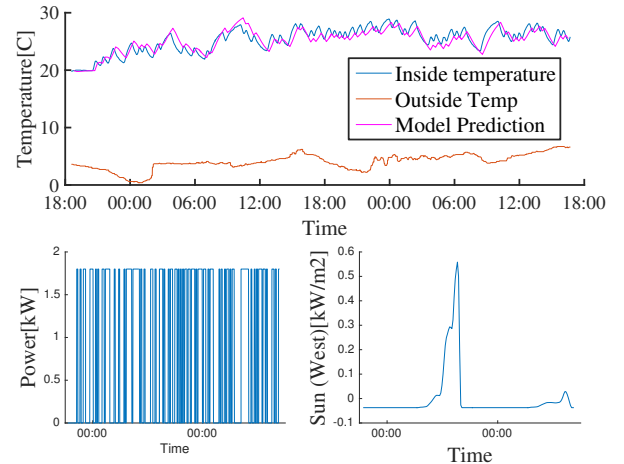


Figure 3: Validation of the model for room NW. The prediction error over this one day period is typically below  $1^\circ\text{C}$ . Upper: Actual measured indoor temperature vs. temperature predicted by the identified model. Outside temperature during the experiment is also displayed. Lower: Remaining inputs driving the system, namely solar radiation (right) and heater input (left). The heater input in this experiment was an open loop PRBS profile

Finally, the individuals models (2) have been transformed from input-output form to state-space form and combined as:

$$\begin{aligned} x_{i+1} &= Ax_i + B_u u_i + B_d d_i \\ y_i &= Cx_i \end{aligned} \quad (3)$$

where  $x$  denotes the state of the system. We recall that  $u = \frac{q}{P_{\max}}$ .

The identified model reveals the main characteristics of the building: the impact of the sun on the rooms is significant: static gains from solar inputs to temperature range from  $3$  to  $11^\circ\text{C}/(\text{kW}/\text{m}^2)$  depending on the rooms, while the effect of outdoor temperature is milder with static gains ranging from  $0.1$  to  $0.2^\circ\text{C}/^\circ\text{C}$ (outside). Two time constants were identified for each room: one fast time constants ranging between  $5$  and  $10$  minutes and one slow time constant ranging between  $1$  and  $2.5$  hours. This can be interpreted as follows: the faster time constant corresponds to the air thermal mass that can be heated directly with our heating system. The fastest variation of temperature in Figure 3 are associated with this part of the dynamics. The slower time constant corresponds to heavier thermal mass of the building. This second mode is the one that the controller will try to utilize to store energy in the system over few hours. Finally, rapid steady state calculations suggests that the equivalent average U-value of the envelope of the building is between  $1$  and  $1.5\text{W}/\text{m}^2/\text{K}$ . This value is not particularly high, especially considering that the construction is  $40$  years old and it is also not very low, meaning that the insulation is not extremely performant. Therefore, this building can be considered as representative of an average office building in Switzerland. Note that the dynamics of the building has a direct influence on the

regulation capacity it can offer. In particular, it should be expected that very efficient buildings will not be able to offer considerable capacity due to very low energy consumption related to heating, cooling and ventilation.

#### 4. Reserve scheduling

The reserve scheduler computes two quantities: a baseline energy consumption for the next day  $\bar{\mathbf{p}}_{|DA}$ , and a capacity bid,  $\gamma$ . Conceptually, the computed capacity bid should be chosen considering the following:

- For a given capacity bid, the controller needs to schedule the baseline consumption in such a way that regardless of the AGC signal it receives, it can shift its power consumption by that amount and still satisfy operational constraints.
- The building operator receives a payment for the “flexibility” he offers which is proportional to the bid. On the other hand, it also pays for baseline power. Intuitively, a higher bid also requires a higher baseline (since the building needs to be able to decrease its consumption by larger amounts). Therefore, there is a financial trade-off between bid and baseline. Our previous work [37] has shown that this trade-off depends on the ratio between the price of energy and the reward for flexibility. If the reward for flexibility is high enough, maximizing the bid we can offer will be optimal. For the sake of demonstration, we assume that it is the case. We refer the reader to [37] for a detailed discussion of the financial reward associated with AGC tracking.

We discuss next how to compute the capacity bid. A planning problem is solved at the reserve scheduler level every day for the next regulation period. Two methods to approximate and solve this problem have been discussed in our previous work [28]. This section discusses the main ingredients of the problem formulation and how it is adapted to include the participation in the intraday market. We refer the reader to [28] for a more detailed discussion of the solution methods.

##### 4.1. Building feasible set

We utilize the model of the system as described in Equation (3) to compute the capacity bid. Knowing the model of the system and the operational constraints, we can formulate the set of all feasible input trajectories that the building can follow while respecting constraints. It takes the following form:

$$\mathcal{U}(\bar{x}, \tilde{\mathbf{d}}) = \left\{ \mathbf{u} \left| \begin{array}{l} x_{i+1} = Ax_i + B_u u_i + B_d d_i \\ y_i = Cx_i \\ |y_i - T_{\text{ref}}| \leq \beta \\ u_i \in \mathbb{U} = [0, 1]^{n_u} \\ x_0 = \bar{x}, \\ \forall i = 0, \dots, N_1 - 1, \end{array} \right. \right\} \quad (4)$$

where  $N_1$  is the horizon covering the participation period,  $T_{\text{ref}}$  the optimal temperature and  $\beta$  a parameter controlling the allowed comfort level deviation from optimum.  $\mathcal{U}$  captures the dynamics of the buildings through the first two equations, starting from the current estimated initial condition  $\bar{x}$ . Equation  $|y_i - T_{\text{ref}}| \leq \beta$  stipulates that the zone temperatures should not deviate from the reference temperature by more than  $\beta$ , a design parameter that we will vary to control the level of comfort. The constraint  $u_i \in \mathbb{U}$  captures the constraint on maximum and minimum PWM ratio to control the power of the heaters.

Notice that this set depends on the initial condition  $\bar{x}$  and a forecast for the disturbance  $\tilde{\mathbf{d}}$ . Finally an expression of the power consumption as a function of the control inputs is required. In our case it takes the simple form  $p = h(u) = uP_{\text{max}}$ , due to the purely resistive nature of the heaters.

##### 4.2. Modelling of the AGC signal

One of the objective of our controller is to ensure that tracking can be achieved for a wide range of potential AGC realizations. To that effect, we build a set  $\Xi$  based on previously observed realizations of the AGC as:

$$\Xi = \left\{ \sum_{j=1}^{N_s} \lambda^{(j)} \mathbf{a}^{(j)} \mid \sum_j \lambda^{(j)} = 1, \lambda^{(j)} \geq 0 \right\} \quad (5)$$

where the  $\mathbf{a}^{(j)}$ 's are previously observed realizations of the uncertainty.

[28] proposes another alternative for the choice of  $\Xi$  and discusses the implications of the choice of  $\Xi$  in detail. The key idea is that if the controller is able to handle values of the AGC that have been observed in the past it should perform well for new realizations due to the consistency of the AGC over time.

##### 4.3. Computation of the bid

We introduce the following robust optimization problem:

#### Problem 1 (Reserve Scheduling Problem).

$$\begin{aligned} & \text{maximize} && \gamma \\ & \text{s.t.} && \\ & \text{(Recourse policies)} && \mathbf{u}(\mathbf{a}) = \pi(\mathbf{a}) \quad \forall \mathbf{a} \in \Xi \quad (6) \\ & && \bar{\mathbf{p}}_{|\cdot|-\delta}(\mathbf{a}) = \kappa(\mathbf{a}) \quad \forall \mathbf{a} \in \Xi \quad (7) \\ & \text{(Building Constraints)} && \mathbf{u}(\mathbf{a}) \in \mathcal{U}(x_0, \tilde{\mathbf{d}}) \quad \forall \mathbf{a} \in \Xi \quad (8) \\ & \text{(Power Consumption)} && \mathbf{p}(\mathbf{a}) = h(\mathbf{u}(\mathbf{a})) \quad \forall \mathbf{a} \in \Xi \quad (9) \\ & \text{(Power tracking)} && \|\epsilon(\mathbf{a})\|_{\infty} \leq \alpha\gamma \quad \forall \mathbf{a} \in \Xi \quad (10) \end{aligned}$$

The decision variables are the capacity bid  $\gamma$ , the day-ahead baseline consumption  $\bar{\mathbf{p}}_{|DA}$ , and the control policies  $\pi$  and  $\kappa$ .  $x_0$  and  $\mathbf{d}$  are data of the problem and represent the initial condition of the system and the prediction for

the disturbances affecting the system, namely the weather and internal gains.

Equation (8) captures the operational constraint of the building and Equation (9) the relationship between control inputs to the system and power consumption, both described in Section 4.1. Equation (10) is the power consumption tracking constraint as discussed in Equation 1.

In this problem, we know that the uncertain tracking request  $\mathbf{a}$  will be revealed progressively over time. Therefore,  $\mathbf{u}$  and  $\mathbf{p}$  do not need to be decided at the bidding stage, but can be functions of the uncertain parameter  $a$  as expressed in Equations (6) and (7) where  $\pi$  and  $\kappa$  are functions and need to respect some causality requirements. The decision variables have to be chosen so that constraints are satisfied for every value of the AGC in  $\Xi$ . Therefore this problem is a robust optimization problem.

In this work, the goal is to demonstrate that we can evaluate successfully the regulation capacity that the system can offer. Therefore we choose to maximize the simple cost function  $\gamma$  so as to maximize the capacity offered. In an economic optimization, we would need to include the cost of the energy purchased. If the payment for capacity is large compared to the cost for energy, the result will be virtually identical while the regulation capacity offered will decrease as the cost of energy increases. The previous work [37] analyzes the economics of frequency control participation in Switzerland and has identified that it is most of the time worthwhile to offer large capacities when possible.

**Remark 2.** The actual disturbance  $\mathbf{d}$  is not known exactly at the time the problem is solved, only a forecast is available. Conceptually, it could be treated exactly like  $\mathbf{a}$  and the solution of the problem could be robustified against forecast error, for which data is readily available (at least for weather). The control decision would then become a function of the disturbance, *i.e.*  $u = \pi(\mathbf{a}, \mathbf{d})$ . In the optimization, we assume the disturbance will take its nominal value. The reason for this are: 1) the variability in the AGC is larger than in the disturbance and is the main focus of this work, 2) forecast are good enough on short timespan (one day), so the gain in explicitly modelling the disturbance uncertainty does not necessarily justify the increased computational cost and modelling effort.

#### 4.4. Constraints on the policies

As such, Problem 1 is infinite dimensional due to the presence of the policies  $\pi$  and  $\kappa$ . It is customary [38] to parametrize these policies as affine policies due to their nice trade-off between performance and computational properties. It results in  $\mathbf{u} = \mathbf{M}\mathbf{a} + \mathbf{v}$  and  $\bar{\mathbf{p}}_{\cdot|i-\delta} = \mathbf{N}\mathbf{a} + \mathbf{w}$ . In order to ensure that the policies are causal, appropriate constraint on  $\mathbf{M}$  and  $\mathbf{N}$  are imposed so that  $u_i$  can depend on  $a_i$  up to time  $i$  and  $\bar{p}_{i|i-\delta}$  on  $a_i$  up to time  $i - \delta$ .

Namely, we impose that:

$$\begin{aligned} M_{i,j} &= 0 \text{ for } j > i \\ N_{i,j} &= 0 \text{ for } j > i - k \end{aligned}$$

where  $k$  denotes the number of time steps per hour in the optimization model, so that the intraday trade cannot depend on observations later than an hour before delivery as per the market rules.

## 5. Closed loop control

This section details the two lower-level layers of the control architecture. Once the capacity bid has been computed together with the baseline, the task of the controller is to satisfy the tracking constraint while making sure that operational constraints are simultaneously met.

### 5.1. MPC controller

We propose to use a predictive controller to maintain comfort satisfactorily. The controller relies on the assumption that being at maximum comfort also maximizes the flexibility. It is the case if the maximum comfort temperature is chosen as the center of the comfort constraint range and uncertainty is symmetric (that is positive and negative AGC are equally likely). The steps of the MPC controller algorithm are:

1. Collect most updated current forecast for weather  $\tilde{\mathbf{d}}$ . This forecast is recovered from different web services. See section 7.1 for more details on the weather prediction.
2. Form a forecast  $\tilde{\mathbf{a}}$  for the average of the AGC over the next few hours sampled at 15 minutes. It has been shown in [29] that the AGC is a time-correlated signal, at least up to two hours ahead. A predictor for the AGC over the prediction horizon is used: it exploits the time-correlation properties of the AGC over short timescales to predict ahead.
3. Solve the following MPC problem

$$\begin{aligned} \text{minimize} \quad & \|\mathbf{y} - T_{\text{ref}}\|^2 \\ \text{s.t.} \quad & \forall i = t, \dots, t + N_2 - 1 \\ & x_{i+1} = Ax_i + B_u u_i + B_d \tilde{\mathbf{d}}_i \\ & y_i = Cx_i \\ & \mathbf{u} \in \mathcal{U}(x_t, \tilde{\mathbf{d}}), \\ & \|\tilde{\epsilon}\|_{\infty} \leq \alpha\gamma \\ & \text{with } \tilde{\epsilon}_i = p_i - \bar{p}_i - \gamma\tilde{a}_i \end{aligned} \tag{11}$$

where the decision variables in this problem are  $\mathbf{u}$ , and  $\bar{p}_{i|i-\delta}$  for  $i \geq t + \delta$ .  $\gamma$ ,  $\bar{\mathbf{p}}$  and  $\bar{p}_{i|i-\delta}$  for  $i < t + \delta$  are fixed in that problem, and come respectively from the reserve scheduling problem and previous iterations of the MPC controller. The controller aims at maintaining the temperature in the middle of the comfort range for the nominal predictions of

the AGC and the disturbance acting on the system. Note that the controller is free to adjust the baseline after a delay of  $\delta$  conformly to the rules of the market.

4. Place an order on the intraday market to buy  $\bar{p}_{t+\delta|t}^*$ , the intraday market trade computed in Problem 11. This effectively adjusts the baseline for timeslot  $t + \delta$ . Pass down the computed control input  $u_t^*$  to the tracking controller. Go back to step 1) at the next iteration.

## 5.2. Tracking controller

The tracking controller receives the AGC signal  $a_t$  at a fast rate and chooses the control input to the radiators. The tracking constraint reads  $\|p_i - \bar{p}_i - \gamma a_i\| \leq \alpha \gamma$ . Upon reception of the optimal control action  $u^*$  from the MPC controller, the lower level controller computes the power input share going to zone  $j$  as  $\nu^j = \frac{u^{*j}}{\|u^*\|}$ , where  $u^j$  denotes the input to zone  $j$ . Using the current value of the AGC, the control input  $u^j = \nu^j(\bar{p} + \gamma a)/P_{\max}$  is computed. This value is capped between 0 and 1 to give the actual input to the heater, which is applied. It is easily seen that this strategy ensure exact tracking as long as the value of  $u^j$  is between 0 and 1. The optimal dispatch  $u^*$  was computed using a forecast for the AGC: as long as the forecast is not widely different from the actual realization, the value of  $u^j$  is close to the optimal value  $u^*$  computed at the upper level. Note that if the forecast of the AGC over the each 15 minutes period was correct at the MPC level, then the temperature prediction was also correct. This is due to the fact that the thermal system is essentially a low-pass filter and fast variation of the AGC will not affect the output of the system. They can therefore be disregarded in the MPC problem without consequence.

## 6. Uncertainty mitigation through intraday trades

A load providing secondary frequency control is subject to large uncertainties of different origin: tracking request of an a priori unknown AGC signal, weather/occupancy/power consumption forecasting errors, unmodeled disturbances (window openings, ...). The goal of the controller is to reject these disturbances. During the bidding problem, the provider decides the magnitude of the AGC signal it will track, and therefore the amount of uncertainty coming from the tracking request it will support. This is explicitly modeled in the bidding problem as described in Section 4. To illustrate the benefit of intraday trades, we study here how they can be used to mitigate the uncertainty in the AGC tracking request. Thanks to the availability of historic data and the simple nature of the AGC, a quantitative study has been carried out.

A secondary frequency control provider needs to increase or decrease its energy consumption in response to an a priori unknown AGC signal. It may happen that

over a period of one day, the AGC is consistently positively or negatively biased, so that the system is required to release or absorb a significant amount of energy. We refer to the integral of the AGC over time as its energy request. AGC signals with the highest energy requests are the most problematic to handle since they require the system to store or release significant amounts of energy. This can cause two types of issues: 1) the limited energy storage capacity of the system might not support such a request, 2) the system operates away from its optimal operating point which might degrade the performance of the equipment. Through the use of intraday trades, it is possible to reset the energy request close to zero by applying an appropriate filtering strategy, as discussed in [29]. The filtering strategy consists in measuring the current energy request of the sum of the AGC and previous trades on the intraday market, then purchase the negative of that quantity on the intraday market for time slot  $t + \delta$ . This strategy attempts to reset the energy request of the AGC to zero at every timestep, but is affected by the delay of one hour caused by the clearing of the market. A ‘filtered’ AGC signal is obtained that way as the sum of the actual AGC and the intraday trades. Figure 4 shows how the energy request histogram is transformed by applying this filtering strategy. This filtering strategy is compared to the one where the state of charge is reset only once per day, as would be the case if the provider only uses the day-ahead market to buy and sell energy. This shows that the energy storage capacity of the physical system needed is greatly reduced if we resort to the intraday market to offset its ‘state of charge’. It is noticeable that the use of the intraday market reduces the variance of the energy request and brings the storage requirement typically around one kWh per kW of capacity offered against three to four times more when doing only day-ahead purchases.

**Remark 3.** Compensating large energy request by trading on the intraday market incurs a cost associated with the intraday transactions. It is shown in [37] that if the intraday market is only used to compensate for large energy requests, the cost incurred is: 1) most of the time small in magnitude compared to the payment for capacity, 2) on average over the year close to zero, 3) compensated in part by the reward for tracking (the deviation in power consumption is paid at a discounted rate if it is positive, reimbursed at a preferential rate if negative). Finally, in closed-loop the system does not systematically have to compensate for deviation, it will only do it so as to bring the operational comfort to acceptable values. That means that the filtering strategy is not explicitly used by the controller, but illustrates the effect the intraday market alone can have in order to mitigate the variability of the AGC.

Similarly, other uncertainty are also mitigated through the use of intraday trades and the opportunity to reschedule the baseline power consumption, limiting the need for very accurate predictions in our approach, as will be discussed in the next section.



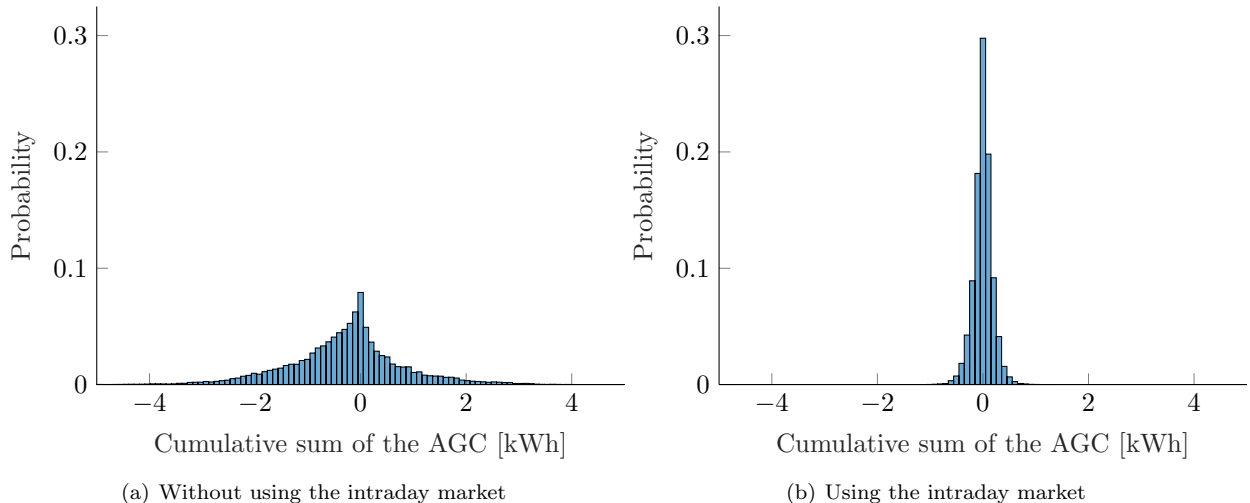


Figure 4: Empirical probability histogram of cumulative sum of the AGC signal over a one-day period (two years of data). The values of the cumulative sum are given as kWh per kW of capacity offered. Left plot shows the histogram for the 'day-ahead filtered' AGC and right plot the 'intraday-filtered' AGC.

## 7. Simulation and Experimental study

### 7.1. Experiments

In this section, we report a series of experiments that has been performed in the period from December 2015 to April 2016.

Each experiment extends over a period comprised between 18 and 24 hours. Experiments differ in two ways:

- Half of the experiment is performed as described in Sections 5 and 4, simulating the building participating in the Swiss regulation market: Proposing a bid for tracking and purchasing energy in the day-ahead and intraday market. In the other half, it is assumed that the building cannot trade energy on the intraday market: therefore its baseline is entirely purchased day ahead and it has no opportunity to reschedule it. The aim is to highlight the effect of the intraday market in the ability of the building to successfully provide regulation. In practice, no intraday trades can be enforced by simply setting variables  $\mathbf{p}_{\cdot, +\delta}$  in the bidding problem to zero a priori.
- Across the experiments, the magnitude of the temperature comfort range is varied with values of  $\pm 0.5$ ,  $\pm 1$ ,  $\pm 1.5$  and  $\pm 2^\circ C$  respectively around the comfort temperature  $T_{\text{ref}} = 23^\circ C$ . This variation allows to explore the trade-off between comfort and the ability of the building to provide regulation services, which translates to a financial reward. Varying the temperature comfort range is a simple way to control the comfort level.


Table 2 compiles the values of the parameters used in the simulation and experiments. The sampling time of the model in Problem 1 and the controller (11) is 15 minutes.

Parameter	Values
$\alpha$	0.05
$T_{\text{ref}}$	$23^\circ C$
$\beta$	0.5; 1; 1.5; $2^\circ C$
$N_1$	96 (1 day)
$N_2$	24 (6 hours)
$N_s$	200


Table 2: Parameters of the simulation and experiments

A selection of representative previously observed AGC signals has been used in the different experiments. Realizations of the AGC with large energy requests have been included since they should illustrate best the influence of the intraday market.

Figure 5 illustrates how the use of the intraday can be beneficial. In Figure 5(a), the baseline was fixed at the beginning of the experiment. At around 7 am, due to errors in forecast and a request for reducing power consumption (negative AGC), the controller ends up in a situation where it has to violate lower temperature constraints for at least one room since its "budget" for power consumption is too low. The controller then takes a few hours to completely recover. Experiment 5(b) was performed in similar conditions but with the possibility to resort to the intraday market trades. Notice that the initial baseline schedule (purple line) is very close to the one for the first experiment, and the AGC test signal is the same in both experiments. It can be seen that from the moment the temperature starts to drop around 6AM the controller anticipates the risk of constraint violation and purchases extra baseline for the upcoming hours, which eventually avoids constraints violation around 10AM, when all temperatures reach the lower value of the constraints. Besides, by explicitly modelling the fact that the baseline could be readjusted, the optimal capacity computed in that case



(a) AGC tracking without participation in the intraday market.



(b) AGC tracking with baseline readjustments in the intraday market.

Figure 5: Two experiments of AGC tracking. The AGC signal is extracted from historical data of 2014. Upper: Power distribution across the four different zones. Baseline energy consumption purchased on the day-ahead market (solid purple line). For comparison, the final baseline after intraday adjustments are also reported (black solid line). Finally, the blue line represents the total power to be tracked by the system. Middle Up: Temperature variation for different zones. Each color corresponds to the measured temperature in each zone. Middle down: AGC signal variation and capacity bid. Lower: Transaction on the intraday market. Positive transactions which correspond to buy extra energy are displayed in green whereas negative transactions are displayed in pink.

was larger by about 20%.

We will next support the claim that what is observed in that particular experiment should be observable on average across all experiments and is characteristic of the difference between the two control schemes. Figure 6 aims at quantifying the trade-off, between capacity offered and comfort. The capacity is reported as the percentage of the total power installed that could be offered as up/down regulation. Comfort is computed using the *ASHRAE likelihood of dissatisfaction (ALD)* (see [39] for more details). Experiments have been grouped in pairs for readability where each pair was performed in similar conditions (close weather, same AGC signal, and comfort constraints). In each pair, the round marker represents the result of the experiment when using intraday trades and the square marker the one without having intraday transactions. The first observation is that, in similar conditions, experiments using the intraday market always resulted in higher capacity bids and higher comfort.

Ultimately, each colour of experiment correspond to a different level of constraint tightening  $\beta$  from  $\pm 0.5^\circ C$  in red to  $\pm 2^\circ C$  in green. As expected, as the constraints are relaxed, higher bids for regulation can be offered, and result on average in lower comfort.

A series of simulations were also performed to compare with the experiments. The weather recorded for 60 different days during winter 2015 was used for simulations. For each weather scenario, the optimal bid is computed for 5 different levels of comfort tightening (the same as in the experiment plus  $\pm 0.25^\circ C$ , depicted in dark blue). The optimal computed controller was applied for 9 representative scenarios of the AGC. The resulting level of comfort was computed and reported in Figure 6. The conclusion derived from the experiments are confirmed and can be summarized as follows:

- Trading energy on the intraday market to readjust the baseline allows to offer higher capacity bids and improve comfort while offering regulation.
- Comfort level and regulation capability can be traded off, for example by relaxing the temperature constraints. The more the relaxation, the lower the comfort and the higher the bid.

Secondly, it can be seen that experimental results are consistent with the simulations in the sense that computed bids are almost identical while comfort levels are close. Experiments (especially when constraints are very tight), tend to display lower comfort with respect to the simulations. This is expected since simulations assume perfect predictions and perfect measurements, which of course is not the case in experiments.

Another conclusion seems to appear through the simulation results: for a given constraint level, the use of the intraday market increases the capacity offer, but also mitigates the variance of the comfort with respect to weather scenarios and AGC signals.

To illustrate the extent to which the intraday market is used, we report in Table 3 the statistics of the experiments. The numbers reported the total Energy consumption, total net and absolute amount of intraday trades, as well as total day-ahead energy purchased. For the simulation columns, the number are average over simulations performed with the same set of experiment as for Figure 6, with an extended test set for the AGC scenarios. The average over the 10 experiments reported in Figure 6 is reported for experiments. The following observations are in order:

- The volume of intraday trades (sum of the absolute value of intraday trades) amounts to about 25% of the total power consumption, both in simulations and experiments. This demonstrates that the closed loop controller described in Section 5.2 is not trying to overact on the intraday market to maintain the temperature in the comfort range.
- In simulations, the net intraday energy trades are quite small on average: this means that intraday trades tend to cancel out over time, leaving a net intraday purchase below 5 % of the total power consumption.
- In experiment, intraday trades are consistently positive at around 10 % of the total power consumption. This means that the algorithm tends to underestimate the needed power consumption slightly. Beside the fact that the number of experiment is not statistically significant, it is extremely complicate to identify a single factor explaining this phenomenon: errors in weather forecast, bias in the AGC signal received, model mismatch and unexpected disturbances will together contribute to these prediction errors. Note that if consistent bias was confirmed over a more extensive set of experiments, it should be possible to eliminate it by, e.g., readjusting the prediction model. However, it should be expected that prediction errors cannot be completely eliminated.
- In experiments, the average amount of constraint violations is almost divided by two when resorting to intraday trades. This confirms the observations made based on Figure 5 and can be explained simply: in the case where the baseline cannot be readjusted, the control authority available after the baseline has been computed is relatively limited, whereas it is significantly increased when intraday trades are available. Therefore in the latter case, the controller is able to reject disturbance more efficiently and therefore mitigate constraint violations and increase average comfort, despite the fact that the regulation capacity offered is even larger. A by-product of that is that the controller is then less sensitive to forecast errors, which supports the decision in that case to not model them as uncertain parameter entering the bidding problem.




Figure 6: Tracking capacity bid versus comfort achieved in closed-loop experiments and simulation . Colors correspond to different levels of constraint tightening (dark blue for  $\pm 0.25^{\circ}C$ , red for  $\pm 0.5^{\circ}C$ , orange for  $\pm 1^{\circ}C$ , light blue for  $\pm 1.5^{\circ}C$ , green for  $\pm 2^{\circ}C$ ). White face markers correspond to simulation instances and colored ones to experiments. Circles correspond to experiments with intraday trades and squares to experiments without intraday trades. Desired improvement direction is to the top left

Quantity (in kWh/h)	Simulation (intraday)	Simulation (no intraday)	Experiment (intraday)	Experiment (no intraday)
Total Consumption	3.2	3.3	4.1	3.8
Day Ahead Purchases	3.3	3.3	3.6	3.8
Net Intraday Purchases	-0.1	0	0.5	0
Absolute Intraday Purchases	1	0	0.8	0
Avg Constraint Violation ( $^{\circ}C$ )	0	0	0.026	0.047

Table 3: Statistics of the experiments. Numbers reported are normalized by the length of the experiments, yielding an average hourly consumption. The last line reports the constraint violation, averaged over rooms and time, in Celsius degrees.

One of the take-away message of this paper is that the strength of our approach is that the closed loop controller is able to successfully overcome prediction errors and unexpected disturbance, that are always present in a real-world application, by readjusting the baseline, therefore increasing the robustness of the scheme significantly.

## 8. Discussion

### 8.1. Practical relevance and relations to other works

This work demonstrates how the inherent storage in a building can be used in order to offer significant flexibility at a controlled level of occupant discomfort. Combining our findings with those of other works considering commercial HVAC systems suggests some direction for practical implementation of regulation with commercial buildings. The works [26], and [40] have shown experimentally that frequency regulation could technically be provided by variable speed drive chillers and fans. However, they limit the frequency band supported by pre-filtering the regulation signal within a frequency band between 30 seconds and 10 minutes. By limiting the impact of the controls on the inside temperature as much as possible, only the inertia in the HVAC system is used, and this inertia is quite small in the absence of a storage system. Effectively, the inertia of the building system itself remains unused. By modeling the thermal dynamics of the building, our work demonstrated that a robust strategy could be used to exploit the inertia in the building heated space successfully and, hence, extend the frequency range of the service.

Interestingly, [2] proposes to modulate zone temperature setpoints to achieve frequency regulation, and validates the idea in a simulation case study, but without explicitly modeling the effect of the strategy on the rooms. An interesting research direction could be to combine both concepts. On one side, the inertia of the HVAC could be exploited to absorb the fastest frequencies, for example by changing the duct pressure setpoint [40]. On the other hand, thermal power demand of the indoor space could be used to absorb slower frequencies. This requires the modelling of the response of the room temperature to changes in thermal power input, which has been demonstrated in our work. In general, a model of the effect of changes in the thermal power demand on the electricity power consumption needs to be found, but for slower timescales only, which mitigates the issue of modelling the interactions of all the components on fast timescales. The difficulty of this task is system dependent.

### 8.2. Need for fast actuation

Realistic HVAC system also have a limit in terms of how fast they can vary power consumption. Previous work on chillers and fans suggest that frequency faster than 30 sec can pose operational issues for the equipment. To improve the quality of tracking, it might be needed to attach to the system a fast storage element such as an electric

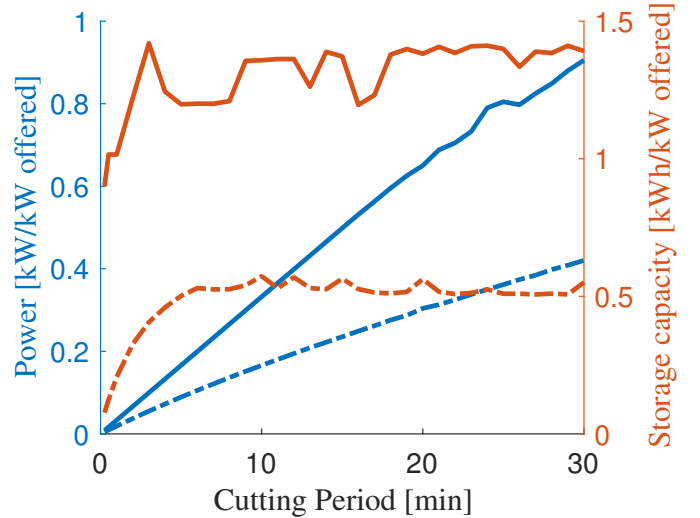


Figure 7: Characteristics of the battery needed to absorb the fastest frequency in the AGC. Power ratings are in blue and energy storage capacity in red. Solid line is the worst case over one year, and dashed line the value needed to cover 99% of the signal. Values are computed for a 1 kW tracking capacity.

battery. Because of the high cost of battery capacity, the operation should try to limit the capacity needed for the control task, and therefore absorb only the fastest frequencies with the battery. Figure 7 gives the size of the battery needed to absorb the high frequency part of the AGC as a function of the filtering frequency.

### 8.3. System benefits

The system used in our experiments offers virtually perfect tracking because of the very fast response of the heating elements and the fact that the tracking requirement was hard coded in the tracking controller (at the expense of occasional temperature violations). This vastly outperforms traditional resources whose tracking response is often subject to ramp constraints, slower response and/or delays. This is one of the benefits of using loads to provide ancillary services. Though the Swiss regulatory framework does not reward financially providers base on their tracking quality, improved tracking improves the network stability. It has actually been recognized and incentivized in other market such as the PJM ancillary services market [2]. On the other hand, the impact of moving part of the regulation power from the transmission grid to distribution networks can be problematic and a system wide study is required to evaluate the consequence of a large deployment of regulation power in distribution networks. Previous studies such as [41] suggest that large amount of frequency control using loads could be deployed without stressing the underlying network, provided an appropriate control architecture is adopted.

## 9. Conclusion

The work presented proposes a model-based control method which computes the amount of grid regulation power a building can offer exploiting its HVAC system and internal thermal storage. Baseline consumption and capacity offered are computed based on a model of the system, which allows to optimally trade comfort and regulation service provision. The scheme proposed successfully models the opportunity to trade on the intraday market to reschedule the baseline. It identifies that higher regulation capacity can then be offered at the same fixed level of comfort. In addition, full-day experiments have been conducted and demonstrate not only that the system can provide regulation power successfully in a realistic setup at a controlled comfort level, but also that the intraday market trades help enhancing the robustness of the control scheme compared to day-ahead baseline scheduling.

Further work will focus on implementing the proposed algorithms on more complex commercial HVAC systems with realistic operational constraints.

## Acknowledgment

We would like to thank Swissgrid Ltd. for providing us with all the relevant data for our simulation and experimental studies.

- [1] D. Callaway, I. Hiskens, Achieving Controllability of Electric Loads, *Proceedings of the IEEE* 99 (1) (2011) 184–199. doi:10.1109/JPROC.2010.2081652.
- [2] P. Zhao, G. P. Henze, S. Plamp, V. J. Cushing, Evaluation of commercial building HVAC systems as frequency regulation providers, *Energy and Buildings* 67 (2013) 225–235. doi:10.1016/j.enbuild.2013.08.031.
- [3] Y. G. Rebours, S. Member, D. S. Kirschen, M. Trotignon, A Survey of Frequency and Voltage Control Ancillary Services Part II : Economic Features, *IEEE Transactions on Power Systems* 22 (1) (2007) 358–366.
- [4] H. Hao, B. M. Sanandaji, K. Poolla, T. L. Vincent, Aggregate Flexibility of Thermostatically Controlled Loads, *IEEE Transactions on Power Systems* 30 (1) (2015) 189–198. doi:10.1109/TPWRS.2014.2328865.
- [5] P. J. Douglass, R. Garcia-Valle, P. Nyeng, J. stergaard, M. Toegeby, Smart Demand for Frequency Regulation: Experimental Results, *IEEE Transactions on Smart Grid* 4 (3) (2013) 1713–1720. doi:10.1109/TSG.2013.2259510.
- [6] S. L. Andersson, A. K. Elofsson, M. D. Galus, L. Gransson, S. Karlsson, F. Johansson, G. Andersson, Plug-in hybrid electric vehicles as regulating power providers: Case studies of Sweden and Germany, *Energy Policy* 38 (6) (2010) 2751–2762. doi:10.1016/j.enpol.2010.01.006.
- [7] F. Oldewurtel, T. Borsche, M. Bucher, P. Fortenbacher, M. G. V. T. Haring, J. L. Mathieu, O. Megel, E. Vrettos, G. Andersson, A framework for and assessment of demand response and energy storage in power systems, in: *Bulk Power System Dynamics and Control - IX Optimization, Security and Control of the Emerging Power Grid (IREP)*, 2013 IREP Symposium, 2013, pp. 1–24. doi:10.1109/IREP.2013.6629419.
- [8] F. Oldewurtel, A. Parisio, C. N. Jones, D. Gyalistras, M. Gwender, V. Stauch, B. Lehmann, M. Morari, Use of model predictive control and weather forecasts for energy efficient building climate control, *Energy and Buildings* 45 (2012) 15–27. doi:10.1016/j.enbuild.2011.09.022.
- [9] Y. Ma, A. Kelman, A. Daly, F. Borrelli, Predictive Control for Energy Efficient Buildings with Thermal Storage: Modeling, Stimulation, and Experiments, *IEEE Control Systems* 32 (1) (2012) 44–64. doi:10.1109/MCS.2011.2172532.
- [10] D. Sturzenegger, D. Gyalistras, M. Morari, R. Smith, Model Predictive Climate Control of a Swiss Office Building: Implementation, Results, and Cost-Benefit Analysis, *IEEE Transactions on Control Systems Technology* PP (99) (2015) 1–1. doi:10.1109/TCST.2015.2415411.
- [11] F. Qureshi, T. Gorecki, C. N. Jones, Model Predictive Control for Market-Based Demand Response Participation, in: *Proceedings of the 19th IFAC World congress*, Vol. 47, Cape Town, South Africa, 2014, pp. 11153–11158.
- [12] H. Hao, A. Kowli, Y. Lin, P. Barooah, S. Meyn, Ancillary Service for the Grid via Control of Commercial Building HVAC Systems, *American Control Conference (ACC)*, (2013) 467–472. doi:10.1109/ACC.2013.6579881.
- [13] H. Hindi, D. Greene, C. Laventall, Coordinating regulation and demand response in electric power grids using multi-rate model predictive control, *Innovative Smart Grid Technologies (ISGT)*, 2011 IEEE PES (January) (2011) 1–8. doi:10.1109/ISGT.2011.5759168.
- [14] T. T. Gorecki, A. Bitlisliolu, G. Stathopoulos, C. N. Jones, Guaranteeing input tracking for constrained systems: Theory and application to demand response, in: *2015 American Control Conference (ACC)*, 2015, pp. 232–237. doi:10.1109/ACC.2015.7170741.
- [15] M. Maasoumy, C. Rosenberg, A. Sangiovanni-Vincentelli, D. S. Callaway, Model predictive control approach to online computation of demand-side flexibility of commercial buildings HVAC systems for Supply Following, in: *Proceedings of the American Control Conference*, 2014, pp. 1082–1089. doi:10.1109/ACC.2014.6858874.
- [16] X. Zhang, M. Kamgarpour, A. Georghiou, P. Goulart, J. Lygeros, Robust optimal control with adjustable uncertainty sets, *Automatica* 75 (2017) 249–259. doi:10.1016/j.automatica.2016.09.016. URL <http://www.sciencedirect.com/science/article/pii/S0005109816303600>
- [17] L. Zhao, H. Hao, W. Zhang, Extracting Flexibility of Heterogeneous Deferrable Loads via Polytopic Projection Approximation, in: *Proceedings of the IEEE Conference on Decision and Control*, Las Vegas, USA, 2016, arXiv: 1609.05966. URL <http://arxiv.org/abs/1609.05966>
- [18] L. Zhao, W. Zhang, A geometric approach to virtual battery modeling of thermostatically controlled loads, in: *2016 American Control Conference (ACC)*, 2016, pp. 1452–1457. doi:10.1109/ACC.2016.7525121.
- [19] A. Nayyar, M. Negrete-Pincetic, K. Poolla, P. Varaiya, Duration-differentiated energy services with a continuum of loads, in: *53rd IEEE Conference on Decision and Control*, 2014, pp. 1714–1719. doi:10.1109/CDC.2014.7039646.
- [20] D. Madjidian, M. Roozbehani, M. A. Dahleh, Battery Capacity of Deferrable Energy Demand, in: *Proceedings of the IEEE Conference on Decision and Control*, 2016, arXiv: 1610.01973. URL <http://arxiv.org/abs/1610.01973>
- [21] A. Bitlislioglu, T. T. Gorecki, C. Jones, Robust tracking commitment with application to demand response, *IEEE Transactions on Automatic Control* (accepted).
- [22] M. Balandat, F. Oldewurtel, M. Chen, C. Tomlin, Contract design for frequency regulation by aggregations of commercial buildings, in: *2014 52nd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2014, pp. 38–45. doi:10.1109/ALLERTON.2014.7028433.
- [23] E. Vrettos, F. Oldewurtel, F. Zhu, Robust Provision of Frequency Reserves by Office Building Aggregations.
- [24] M. Maasoumy, Flexibility of Commercial Building HVAC Fan as Ancillary Service for Smart Grid, in: *IEEE Green Energy and Systems Conference (IGESC 2013)*, 2013.
- [25] Y. Lin, P. Barooah, S. Meyn, T. Middelkoop, Experimental Evaluation of Frequency Regulation from Commercial Build-

- ing HVAC system, *IEEE Transaction on Smart Grid* 6 (no. 2) (March 2015) 776–783.
- [26] L. Su, L. K. Norford, Demonstration of HVAC chiller control for power grid frequency regulation Part 2: Discussion of results and considerations for broader deployment, *Science and Technology for the Built Environment* 21 (8) (2015) 1143–1153. doi:10.1080/23744731.2015.1072455.
- [27] E. Vrettos, E. C. Kara, J. MacDonald, G. Andersson, D. S. Callaway, Experimental Demonstration of Frequency Regulation by Commercial Buildings - Part II: Results and Performance Evaluation, arXiv:1605.05558 [math]ArXiv: 1605.05558. URL <http://arxiv.org/abs/1605.05558>
- [28] L. Fabietti, T. T. Gorecki, F. A. Qureshi, A. Bitlislioglu, I. Lympelopoulou, C. N. Jones, Experimental implementation of frequency regulation services using commercial buildings, *IEEE Transactions on Smart Grid*.
- [29] I. Lympelopoulou, F. A. Qureshi, T. Nghiem, A. A. Khatir, C. N. Jones, Providing Ancillary Service with Commercial Buildings: The Swiss Perspective, in: 9th IFAC International Symposium on Advanced Control of Chemical Processes (AD-CHEM), Whistler, BC, Canada, 2015.
- [30] S. SA, Principes des produits de services système, Tech. rep. (2015). URL [https://www.swissgrid.ch/dam/swissgrid/experts/ancillary\\_services/Dokumente/D150201\\_AS-Products\\_V8R0\\_FR.pdf](https://www.swissgrid.ch/dam/swissgrid/experts/ancillary_services/Dokumente/D150201_AS-Products_V8R0_FR.pdf)
- [31] Swissgrid, Test for secondary control capability, Tech. rep. (2012). URL [https://www.swissgrid.ch/dam/swissgrid/experts/ancillary\\_services/prequalification/D130422\\_Test-for-secondary-control-capability\\_V2R1\\_EN.pdf](https://www.swissgrid.ch/dam/swissgrid/experts/ancillary_services/prequalification/D130422_Test-for-secondary-control-capability_V2R1_EN.pdf)
- [32] L. Su, L. K. Norford, Demonstration of HVAC chiller control for power grid frequency regulation Part 1: Controller development and experimental results, *Science and Technology for the Built Environment* 21 (8) (2015) 1134–1142. doi:10.1080/23744731.2015.1072449.
- [33] Office fédéral de l'énergie (OFEN), Ex-post-analyse des schweizerischen energieverbrauchs 2000 bis 2013, Tech. rep. (sep 2014).
- [34] L. Ljung, System identification, in: *Signal Analysis and Prediction*, Springer, 1998, pp. 163–173.
- [35] U. DOE, EnergyPlus Engineering Reference, 2012.
- [36] Q. Zhang, J. Huang, S. Lang, Development of typical year weather data for Chinese locations/Discussion, *ASHRAE transactions* 108 (2002) 1063.
- [37] F. A. Qureshi, I. Lympelopoulou, A. Ahmadi Khatir, C. Jones, Economic advantages of office buildings providing ancillary services with intraday participation, Tech. rep. (2016).
- [38] Ben-Tal, A., El Ghaoui, L., Nemirovski, A., *Robust Optimization*, Princeton University Press, 2009.
- [39] S. Carlucci, *Thermal comfort assessment of buildings*, Springer, 2013.
- [40] Y. Lin, P. Barooah, S. Meyn, T. Middelkoop, Experimental Evaluation of Frequency Regulation From Commercial Building HVAC Systems, *IEEE Transactions on Smart Grid* 6 (2) (2015) 776–783. doi:10.1109/TSG.2014.2381596.
- [41] M. D. Galus, S. Koch, G. Andersson, Provision of Load Frequency Control by PHEVs, Controllable Loads, and a Cogeneration Unit, *IEEE Transactions on Industrial Electronics* 58 (10) (2011) 4568–4582. doi:10.1109/TIE.2011.2107715.