

Experimental Implementation of Frequency Regulation Services Using Commercial Buildings

Luca Fabietti, Tomasz T. Gorecki, Faran A. Qureshi, Altuğ Bitlislioglu, Ioannis Lympelopoulos, Colin N. Jones

Abstract—This paper illustrates the potential of commercial buildings to act as frequency reserves providers through an experimental demonstration conducted in a multi-zone university building. The proposed control methodology is presented in detail, including the control architecture, the controller design, model identification, and hardware description. Finally, the effectiveness of the presented approach is tested by means of simulations and experiments in a controlled environment.

Index Terms—Ancillary services, commercial buildings, area control signal, frequency regulation, model predictive control (MPC).

I. INTRODUCTION

THE need to modify the current structure of the electric grid into a more sustainable configuration motivates a rapid increase in renewable generation. However, renewable energy sources are intrinsically volatile and uncertain, which poses new engineering challenges in order to guarantee the proper and safe functioning of the grid [1].

One of the main instruments for regulating, at any time, the balance between generated and consumed power is a set of reserves also known as ancillary services. Ancillary services can be divided into two categories: the first type corresponds to those services that are activated to respond to normal fluctuations of the grid. The second type is instead used in case of a contingency as, e.g. the loss of a power plant [2].

In this paper, we focus on frequency regulation which falls into the first category. In the Swiss market, the grid operator buys reserve power over a specified period in the future by means of weekly/daily auctions from a set of Ancillary Services Providers (ASPs). During real-time operation, the grid operator calculates the mismatch between generation and production and computes the so-called Automatic Generation Control Signal (AGC signal) required to close this gap. The AGC is transmitted in real time to all participating ASPs that activate according to the size of their accepted bids in the auction.

Historically, these kinds of services were offered by the supply-side which adapts power generation according to the

AGC signal. However, maintaining generation resources in standby comes at the cost of increased operating and monitoring expenses. Moreover, the AGC is sampled at a high frequency and even highly responsive generators are, in general, not capable of offering precise power tracking [2].

Recent research has pointed out that the demand-side could provide cheaper and more reliable ancillary services [3], [4]. In fact, because the AGC is close to a zero-energy signal, storage technologies have been widely acknowledged to be the ideal suppliers for this type of service. However, storage using electric batteries has several drawbacks that currently prevent their extensive use: they are expensive, not environmentally-friendly and have a limited functional life. Therefore, the research community has also considered different kinds of controllable loads, such as buildings.

With 40% of the total energy consumption worldwide [5], buildings represent a tremendous untapped energy storage source. Furthermore, buildings are inherently characterized by a large thermal inertia that allows to think of them as virtual storage devices capable of decoupling production from consumption.

The fact that commercial buildings are to be preferred, for providing services to the grid, over residential buildings has been underlined in many contributions [6], [7], [4]. The reasons underpinning this choice can be summarized in the following: 1) They are typically characterized by a larger HVAC system with respect to residential buildings. This corresponds to higher energy consumption which, due to the current structure of the energy market, makes it easier for them to be recognized as providers of ancillary services 2) The power consumption of most commercial buildings can be varied continuously as opposed to typical switched control systems present in residential buildings 3) Most commercial buildings are already equipped with energy management systems which could facilitate the communication with the grid operator and allow simpler variation of their energy consumption.

This paper presents experiments of AGC tracking carried out over 10 hours periods in a 90 sq. meters ($\sim 1,000$ sq. ft.) section of the Automatic control Lab at École Polytechnique Fédérale de Lausanne (EPFL) with four different zones. Two model-based methods are proposed to compute the optimal flexibility bid, compared in simulation, and implemented on the real system.

A. Related work

Model-based control strategies such as Model Predictive Control (MPC) have been extensively studied and successfully

The authors are with the Automatic Control Laboratory in EPFL, Lausanne, Switzerland.

{luca.fabietti, tomasz.gorecki, faran.qureshi, altug.bitlislioglu, ioannis.lympelopoulos, colin.jones} @epfl.ch

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.

applied on buildings in order to reduce their operating costs [8], [9]. Several simulation studies have also been conducted on buildings showing their capability of offering ancillary services. In [10], *Hao et al.* propose a control architecture to track the regulation signal by modulating power consumption of the ventilation system. In another interesting contribution, *Hindi et al.* [11] design a multirate MPC controller to jointly manage the provision of both regulation and demand-response to compensate for demand-supply imbalances. In [12] the authors describe a formal method to characterize, as a virtual battery, the deviation in power consumption that a building can support around a nominal consumption. Another relevant contribution is provided by [13] where *Maasoumy et al.* propose an MPC min-max approach to define and quantify the admissible flexibility of a single-zone commercial building. Finally, aggregation of multiple buildings is considered in [14] and [15], where *Vrettos et al.* propose a three-level control architecture for the provision of frequency control reserves by a set of office buildings.

All the previously mentioned works focused on simulations; on the experimental side, very few contributions have been published. In [6], [13] authors empirically estimate the potential power flexibility of a university building. The power consumption of the main ventilation system is indirectly varied acting on the supply duct static pressure setpoint value. Different experiments are conducted showing how the consequent fluctuating air mass flow does not affect occupants' comfort. Another published example of implemented HVAC control for frequency regulation is [7] where power tracking is provided by adjusting fan power consumption. In particular, the building receives a filtered version of the AGC signal which is then tracked with respect to the nominal baseline computed by the pre-existing building controller. Experiments were conducted on a 40,000 sq. ft. over a 40 minutes time duration.

B. Contribution of the paper

Based on our previous works [12], [16], in this paper we present the design of two MPC based controllers for the provision of frequency regulation services using multi-zone commercial buildings. We extend previous results, firstly providing a formal method to quantify the flexibility of a commercial building in a multi-input framework. The construction of the uncertainty sets based on real recorded data is discussed, and two different strategies to solve the problem are compared with extensive simulations. Secondly, at-scale experiments of AGC tracking following the rules imposed by the current regulation of the electricity market in Switzerland are reported. To the best of the authors' knowledge, this work represents the first report of experimental results for frequency regulation, based on real recorded AGC signals, using a commercial building over a realistic time range, and the first work that applies a formal method to compute an optimal bid for AGC provision on an real system.

C. Structure of the manuscript

The rest of this paper is organized as follows. In Section II we present the problem formulation of ancillary service provision using commercial buildings. In Section III, we describe

TABLE I
PARAMETERS DESCRIPTION

Symbol	Description
a	AGC signal
γ	Capacity Bid
\bar{p}	Baseline Electric Power Consumption
$p_{\text{ref}} = \bar{p} + \gamma a$	Power Consumption Tracking Reference
$\epsilon = p_{\text{ref}} - p$	Tracking error
u	Command to the HVAC
q	Thermal input to the building
p	Power consumption of the HVAC
y	Output of the building
w	Weather Disturbance

the controller design. In Section IV, simulation studies are analyzed. Section V describes experimental results. Section VI provides a discussion on the applicability of the method proposed, while Section VII summarizes our conclusions and proposes some future improvements and directions.

II. PROBLEM STATEMENT

In the following, the formulation of the problem is introduced and tailored to the particular structure of the Swiss electricity market. In most countries, the process of frequency ancillary services is, however, similar. The reader is referred to [17] and reference therein for an exhaustive overview of frequency and voltage control ancillary services in power systems.

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., $\mathbf{p} = [p_0^T, p_1^T, \dots, p_{N-1}^T]^T$.

The tracking service procurement, for the case of buildings, involves two phases as illustrated in Figure 1: bidding and tracking.

The bidding phase is done offline and consists in advertising to the grid operator two quantities, namely a *baseline* energy consumption \bar{p} , and a *capacity bid* γ . The latter represents the highest deviation (in absolute value) in power consumption with respect to the purchased baseline the ASP is willing to track over the activation period. More specifically, if a bid of $\gamma = 1MW$ is accepted, the provider will then receive an AGC signal with a maximum magnitude of $\pm 1MW$.

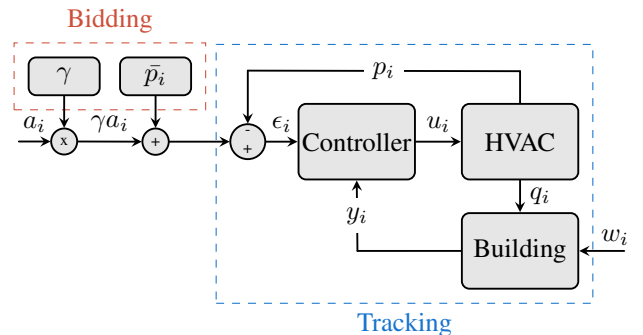


Fig. 1. Architecture of the control system for tracking service procurement.

The second phase is tracking. This phase is performed online as follows: considering the current state of the network, the power system operator will compute, at time i , a

normalized AGC signal a_i , which will be sent to the service providers according to their respective accepted bids, so that the magnitude of the AGC signal they receive is proportional to their bid. The providers then have to modify their power consumption p_i so that

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

where \bar{p}_i is the baseline purchased for timestep i , p_i the actual power consumption at timestep i , a_i the normalized AGC signal and γ is the capacity bid. r represents the maximum tracking error allowed, as a percentage of the total submitted bid. This essentially states that the provider must consume the sum of the baseline power it has purchased and the scaled AGC signal it receives and that its tracking error ϵ_i must be less than a percentage of the capacity bid γ .

The goal of the controller is to choose the inputs to the HVAC system u in order to maintain the tracking error small enough as specified by equation (1), and, at the same time, satisfy constraints on the system.

At the time of bidding, the provider is therefore faced with an economic decision that can be conceptualized by the following optimization problem:

Problem 1 (Bidding Problem).

$$\begin{aligned} & \text{minimize} && \mathbb{E}_{\mathbf{a}}[J(\gamma, \bar{\mathbf{p}}, \mathbf{a})] \\ & \text{s.t.} && \forall \mathbf{a} \in \Xi \\ & && \mathbf{u} \in \mathcal{U}(x, \psi, \mathbf{w}), \\ & && \|\epsilon\|_{\infty} \leq q\gamma \end{aligned} \quad (2)$$

where the tracking signal \mathbf{a} is uncertain and the decision variables are the capacity bid γ , the committed baseline consumption $\bar{\mathbf{p}}$, the input trajectory \mathbf{u} and the power consumption trajectory \mathbf{p} . $\mathcal{U}(x, \psi, \mathbf{w})$ is a set representing the admissible input trajectories for the system and whose construction is detailed in Section II-B.

At the time of decision, the AGC signal \mathbf{a} is not known yet. Therefore this problem is an ‘‘uncertain’’ optimization problem [18], [19]. The cost function J captures the total cost of operation and depends on the power consumption \mathbf{p} , and the AGC signal \mathbf{a} . Since the uncertainty is not yet revealed, it is minimized in expected value. Constraints, on the other hand, are handled in a robust fashion, and need to be satisfied *for all values of the uncertainty* in a set Ξ . This modeling choice is motivated by the market rules: the tracking requirement holds at all times irrespective of the tracking signal a .

The AGC signal a is revealed progressively and it is possible to re-adjust the control action u accordingly, therefore the goal of Problem 1 is not to fix decisions \mathbf{u} once and for all, but to rather optimize over control policies so that $\mathbf{u} = \pi(\mathbf{a})$. This type of uncertain optimization is referred to as a multi-stage optimization problem. In the bidding problem, the first stage variables are the capacity γ and the baseline $\bar{\mathbf{p}}$, while variables \mathbf{u} and \mathbf{p} are later stages variables. The problem therefore has $N+1$ stages. Notice that the total duration of the bid is fixed in advance (for example one day), and is divided into N stages, according to how often the uncertainty is revealed and the control decision readjusted. Problem 1 can be thought of as an optimization over causal control policies π , *i.e.* decisions

at stage k depend only on observation of the uncertainty up to time k .

In order to solve Problem 1, the following elements need to be defined:

- the choice of the cost function J (Section II-A)
- the set of feasible trajectories \mathcal{U} (Section II-B)
- the choice of the uncertainty set Ξ (Section II-C)
- the parametrization of the control policy π (Section III)

A. Choice of the cost function J

The building receives a payment proportional to the capacity bid and pays for the energy it purchases. This part of the cost is deterministic. The balance between the two depends on the difference between the cost of power and the unit reward price for capacity. In this work, the goal is to illustrate AGC tracking by buildings, therefore we maximize the capacity participation, *i.e.* $J = -\gamma$. It has been shown in a previous work [16] that most of the economic benefit in participating in ancillary services provision comes from the capacity bid. Therefore this is reasonable simplification. We refer the reader to [16] for the treatment of the full cost function.

B. Building and HVAC modeling

A dynamic model of the building and the HVAC is necessary. We assume that the models take a discrete-time state-space form:

$$\text{HVAC:} \quad \begin{cases} \psi_{i+1} &= f_{\text{HVAC}}(\psi_i, u_i, w_i) \\ q_i &= g(\psi_i, u_i, w_i) \\ p_i &= h(\psi_i, u_i, w_i) \end{cases} \quad (3)$$

$$\text{Building:} \quad \begin{cases} x_{i+1} &= Ax_i + B_q q_i + B_w w_i \\ y_i &= Cx_i \end{cases} \quad (4)$$

where x_i and ψ_i are the state vectors describing the building and HVAC system, respectively, u_i the command inputs to the HVAC system (set-points, switching on/off sequences, etc.), q_i the vector of thermal power inputs provided by the HVAC to each zone, p_i the electric power consumption of the HVAC, y_i the zone air temperatures at time step i , and w_i the vector of external disturbance affecting the system. The system is subject to a number of constraints. We define comfort constraints at the level β as $|y_i - T_{\text{ref}}| \leq \beta$. The inputs are also constrained to lie in a convex set which describes physical limitations for the equipment $u \in \mathbb{U}$.

We define the *admissible input trajectories* set as:

$$\mathcal{U}(\bar{x}, \bar{\psi}, \mathbf{w}) = \left\{ \mathbf{u} \left| \begin{array}{l} \psi_{i+1} = f_{\text{HVAC}}(\psi_i, u_i, w_i) \\ q_i = g(\psi_i, u_i, w_i) \\ x_{i+1} = Ax_i + B_q q_i + B_d w_i \\ y_i = Cx_i \\ |y_i - T_{\text{ref}}| \leq \beta \\ u_i \in \mathbb{U} \\ x_0 = \bar{x}, \quad \psi_0 = \bar{\psi} \\ \forall i = 0, \dots, N-1, \end{array} \right. \right\} \quad (5)$$

where \bar{x} , and $\bar{\psi}$ describe the current state of the system. $\mathcal{U}(\bar{x}, \bar{\psi}, \mathbf{w})$ represents the set of all the input trajectories that

preserve occupants comfort while respecting physical limits of the actuators. Note that the building is also affected by external perturbations w including weather and internal gains. It is assumed that these perturbations can be forecast. In the bidding phase, they are, therefore, assumed to be known in advance and we enforce $\mathbf{u} \in \mathcal{U}(\bar{x}, \bar{\psi}, \bar{\mathbf{w}})$ with $\bar{\mathbf{w}}$ the forecast for the disturbance. We will see in Section V-C how experiments demonstrated that forecasting errors did not compromise the performance of the algorithm in a significant way.

C. Uncertainty set construction

In this section, two methods for the construction of the uncertainty set related to the AGC signal are presented.

In the first method, the uncertainty set is obtained by constructing the convex hull of a finite number, N_s , of past realizations of the AGC signal a . The analytic description of the resulting set, Ξ_{ts} , reads:

$$\Xi_{ts} = \left\{ \sum_{j=1}^{N_s} \lambda^{(j)} \mathbf{a}^{(j)} \mid \sum_j \lambda^{(j)} = 1, \lambda^{(j)} \geq 0 \right\} \quad (6)$$

For the second method, a data-driven design approach is also used to construct the uncertainty set. It is known *a priori* that the normalized AGC signal has magnitude at most one. Therefore, we have $\Xi \subseteq \mathcal{B}_\infty(1) = \{\mathbf{a} \mid \|\mathbf{a}\|_\infty \leq 1\}$. Using the uncertainty set $\Xi = \mathcal{B}_\infty(1)$ disregards statistical information about the uncertainty available under the form of scenarios. An approach inspired by [20] can be used for uncertainty set design. The key idea is to use the scenarios to fit the uncertainty set in an optimization problem. Focusing on the integral of the AGC signal, we notice that: a) it is not too large in past scenarios, b) it intuitively is a key factor in determining if the system can support this signal. We propose to build the set Ξ as follows:

$$\Xi_{ms} = \{\mathbf{a} \mid \|\mathbf{a}\|_\infty \leq 1, \|\text{cumsum}(\mathbf{a})\|_\infty \leq s_{\max}\} \quad (7)$$

where $\text{cumsum}(\mathbf{a})$ denotes the cumulative sum of the signal \mathbf{a} ($(\text{cumsum}(\mathbf{a}))_i = \sum_{k=1}^i a_k$). s_{\max} is chosen as the worst-case cumulative sum on a number N_s of previously observed AGC realizations.

A conceptual sketch of the two uncertainty sets is given in Figure 2. In addition to the historical scenarios, Ξ_{ms} contains more extreme realizations that are less likely to be observed in reality. If the same set of samples is used in both approximations, $\Xi_{ts} \subset \Xi_{ms}$ and using the latter set will lead to more conservative solutions.

III. SOLUTION TO THE BIDDING PROBLEM

In this section, we propose two different methods, detailed respectively in Section III-A and Section III-B, to approximate the uncertain bidding Problem 1. Finally, Section III-C is devoted to the description of the online tracking controller.

A. Two-stage approximation

This solution method approximates the bidding Problem 1 by a two-stage robust optimization problem. Instead of having a multi-stage control policy parametrized by $\mathbf{u} = \pi(\mathbf{a})$, a two-stage control policy is assumed. The causality requirements are relaxed and it is assumed that the uncertainty is revealed for the whole prediction horizon at once. First stage variables are the capacity γ and the baseline power consumption $\bar{\mathbf{p}}$ which are decided before the realization of the random parameter \mathbf{a} , while the second-stage decisions \mathbf{u} can depend on the realization of the uncertain parameter. Ξ_{ts} is used as an uncertainty set in this solution method.

For each scenario of the uncertain parameter, there is a separate trajectory of second-stage decision variables resulting in the following optimization problem:

$$\begin{aligned} & \text{minimize} && -\gamma \\ & \text{s.t.} && \mathbf{u}^{(j)} \in \mathcal{U}(\bar{x}, \bar{\psi}, \bar{\mathbf{w}}), \\ & && \|\boldsymbol{\epsilon}^{(j)}\|_\infty \leq q\gamma \\ & && \forall j = 1, \dots, N_s \end{aligned} \quad (8)$$

where the superscript j defines the optimization variables corresponding to the j^{th} scenario of the uncertain parameter and N_s is the total number of scenarios. The baseline power consumption $\bar{\mathbf{p}}$ and the capacity γ are fixed over the whole activation period.

B. Multi-stage approximation

The second approximation scheme we present relies on the main assumption that the system to be controlled is linear. This assumption is discussed in Section VI. The main advantage of this second approach is that the multi-stage structure of the original Problem 1 is preserved. It is known that solving a multi-stage problem based on a scenario tree approach is subject to the so-called ‘‘curse of dimensionality’’, meaning that complexity increases exponentially with the number of stages. A popular approach is therefore to use *decision rules* [21] to parametrize the policy space and reduce it to a finite dimensional one. We use linear decision rules due to their simple form and nice computational properties:

$$u_i = \sum_{j=0}^i M_{i,j} a_j + v_i \quad (9)$$

where $M_{i,j}$ and v_i are the new decision variables. To shorten notation, we vectorize equation (9) as:

$$\mathbf{u} = \mathbf{M}\mathbf{a} + \mathbf{v} \quad (10)$$

The uncertainty set considered in this solution method is Ξ_{ms} . The problem takes the form:

$$\begin{aligned} & \text{minimize} && -\gamma \\ & \text{s.t.} && \mathbf{u} \in \mathcal{U}(\bar{x}, \bar{\psi}, \bar{\mathbf{w}}) \\ & && \mathbf{u} = \mathbf{M}\mathbf{a} + \mathbf{v} \\ & && \|\boldsymbol{\epsilon}\|_\infty \leq q\gamma \\ & && \forall \mathbf{a} \in \Xi_{ms} \end{aligned} \quad (11)$$

Under technical assumption (mostly convexity of the set Ξ_{ms}), problem (11) can be transformed into a convex program. For more details the reader is referred to [12] where the




Fig. 2. Conceptual sketch of the construction of sets Ξ_{ts} and Ξ_{ms} . The grey dots, that represent possible realization of the the unknown AGC signal, are used to fit the two uncertainty sets as described in section II-C.

challenge of optimizing simultaneously the uncertainty set and the control policy is tackled.

Remark 1. *The two approaches differ in their level of conservatism. Problem (11) requires the solution to be immune to all possible signals drawn from the polytope Ξ_{ms} . Secondly, the causality requirement is relaxed in Problem (8), which is an optimistic assumption in the sense that it improves the optimal value of the problem. It can be shown that the optimal objective value function is lower in Problem (8) than in Problem (11). On the other hand, the multi-stage approach has the advantage of retaining the original multi-stage structure and the causality requirement with a limited sacrifice of performance.*

C. On-line operation

In the preceding sections, the bidding problem was detailed. This problem is solved at the beginning of the activation period to determine the capacity bid γ and the baseline \bar{p} . These quantities are then fixed for the duration of the activation. In closed-loop, slightly modified versions of problems (8) and (11) are respectively solved at each time step with a shrinking horizon, and with the baseline power consumption \bar{p} and the power capacity γ being fixed. In the two-stage case, the problem is further modified such that the first time step of the input and the state trajectories are first stage variables while variables for the rest of the horizon are still second stage decision variables. In both cases, the aim of the closed-loop controller is to determine a feasible command input to the HVAC system, u_0 , to each zone so that a high quality tracking service is provided while making sure that thermal constraints are satisfied.

IV. SIMULATION RESULTS

This section compares in simulation the two-stage and the multi-stage approximations discussed in Section III-A, and III-B, respectively.

A. Simulation setup

The model of the experimental setup, described in Section V-B, is used for simulations. The sampling period is chosen equal to 15 minutes which provides a nice compromise between temporal resolution of the control and computational complexity of the problem formulation. The comfort range for

temperature is chosen as 21°C to 25°C ($\beta = 2^\circ\text{C}$). Perfect tracking of the unknown AGC signal is also required ($r = 0$).

To solve the bidding Problem 1 and compute the maximum power capacity that the building can support over the activation period, we use both the two-stage and the multi-stage approach. To this end, the two problems (8) and (11) are solved using, as the uncertainty set, the two robust sets Ξ_{ms} and Ξ_{ts} , respectively. In particular, the two sets are constructed offline as described in Section II-C. The scenarios used to perform this task are obtained by breaking the yearly normalized AGC signal of 2013 into 876 ten hour samples. It is assumed in simulation that the temperature is close to 5°C , which was the one used for steady-state computation. Solving the two bidding problems results in different values for the optimal bid γ^* and baseline \bar{p}^* . To test the robustness, and quality of the solution, historical realizations of the AGC signal of 2014 are considered for validation. The AGC for 2014 is also broken into 876 ten hours test instances. Each ten-hours test AGC sample is multiplied by the optimal power capacity γ and added to the baseline \bar{p} to obtain the total power signal to be tracked by the system. For the two-stage approximation, an open loop optimization problem is solved for each ten hour test sample to optimally distribute the power across the four zones while respecting the comfort constraints. Similarly, for the multi-stage approximation, the optimal affine control law parametrized by \mathbf{M}^* and \mathbf{d}^* is used to compute the open-loop trajectories of the zones temperature. The result is depicted in Figure 3.

B. Analysis of results

As seen in Figure 3 there are a few differences between the two approaches both in terms of bid capacity and of thermal response of the system while providing AGC tracking. The multi-stage approach is more conservative and results in a capacity bid of $\pm 1.85\text{kW}$ while the two-stage approach results in a capacity bid of $\pm 3.2\text{kW}$. This is visible in the bottom plots of Figure 3 where the AGC signals and their maximum amplitude are shown. The computed capacity represents 25% and 43% of the maximum available power, respectively. The resulting temperature trajectories in the four zones of the building for all the considered AGC test samples are shown in the top plots of Figure 3. For the multi-stage approach, the zone temperatures stay more closely around 23°C which represents the most robust state to be in to absorb both positive and negative realizations of the AGC. For the two stage approach, temperatures are closer to the constraints and violate the constraint slightly for a few AGC test samples.

V. EXPERIMENTAL RESULTS

A. Testbed description

The experiments were conducted in four office rooms accounting for a total area of 90 sq. meters. The rooms are characterized by a concrete heavyweight structure with limited glass surface. For the sake of experiments, the original heating system consisting of hot water based radiators was switched off and each room was equipped with a commercial

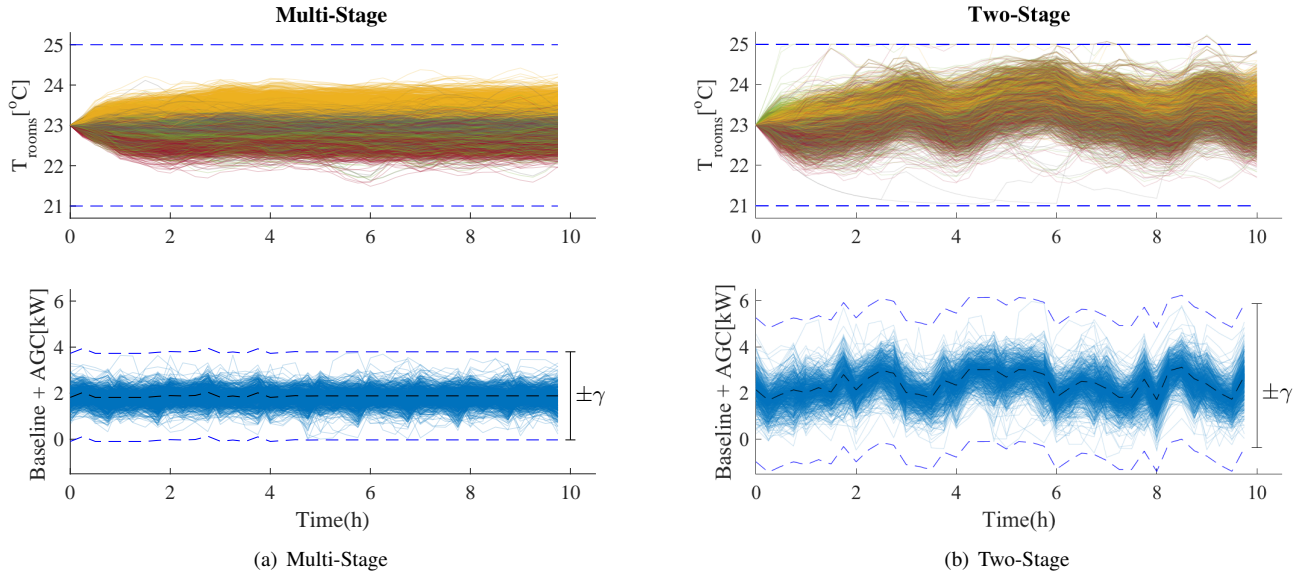


Fig. 3. Open-loop predictions for the AGC multi-stage and two-stage controller for real AGC signals from 2014. The obtained capacity and baseline obtained off-line on data of 2013 are used to compare the two approaches applying 876 different signals extracted from the data set of 2014. Upper: Temperature in each of the 4 zones (different colors) during activation period. Lower: AGC signals superimposed on the computed baseline (black dashed line) and capacity bid bounds.

electric heater. For temperature measurements, Aeotec Z-Wave Multisensor were used. The sensors communicate their measurements to a local server through a Z-Wave network. Finally, all the acquired data are uploaded to a web database with a five minutes frequency.

The heaters are rated at 1850 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.

Power modulation of the heaters is managed by custom BeagleBone Black (BBB) micro-computers. The communications is handled using YARP [?] which is an open-source software supporting cross-platform data interchange. All four heaters are centrally controlled by the tracking controller running on the main server. 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. A schematic of the system is depicted in Figure 4.

B. Modeling and Validation

The model of the HVAC system in our experimental setup is relatively simple. The control input to the system is represented by the pulse-width modulation ratio α to each heater, so that $u_i = \alpha_i$ and $\mathbb{U} = [0, 1]^4$. This directly results in an electric power consumption which is a linear map of the control input, i.e., $p = \alpha P_{\max}$. Finally, the heaters being resistive elements, the power consumption directly translates into a thermal input to the building, so that $q = p = \alpha P_{\max}$.

The model of the building (4) takes as inputs the thermal fluxes from the radiators and outputs the temperatures in each zone. It captures how the heat is circulated in the room.

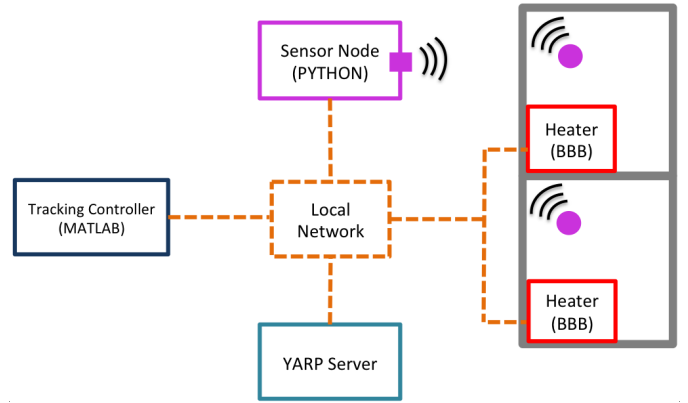


Fig. 4. Schematic of the experimental setup

Each zone was modeled with a second order Auto Regressive model with exogenous inputs (ARX) [22] calibrated using experimental data that were collected during December 2014. The model takes the form:

$$\delta y_{k+1} = a(z)\delta y_k + b_q(z)\delta q_k + b_w(z)\delta w_k \quad (12)$$

where z denotes the delay operator and $\delta y = y - y_{ss}$ with y_{ss} the steady state around which the model is identified. Definitions of δq and δw are identical. Due to nighttime experiments, the main source of disturbance is the outside temperature. Due to their slowly varying nature, no dynamic model was identified for the effect of w : instead, a steady-state relation between q_{ss} and w_{ss} is identified. Through regulation experiments, it has been identified that to maintain an indoor temperature of 23°C when the outside temperature is at an average 5°C , the thermal input required was $q_{ss} = p_{ss} = [0.18, 0.27, 0.72, 0.9]$ kW which corresponds to a command input to the heaters equal to $u_{ss} = [0.10, 0.15, 0.40, 0.50]$. The full model of the building is obtained by combining the individual models of the rooms and has four inputs (heat




Fig. 5. Model validation for one of the 4 considered zones. The identified linear model is used to obtain open-loop predictions over a horizon of 10 hours. The computed trajectory is then compared with the actual measurements.

input in each room) and four outputs (temperatures). It is transformed from input-output form to state-space form as required in equation (4). Figure 5 validates the model for one of the rooms by comparing the actual and the simulated indoor temperature driven by the same input sequence for a time horizon of 10 hours. The model adequately captures the dynamics of the real system with a maximum prediction errors that is smaller than 0.5°C for the whole prediction horizon. Similar results were obtained for the other rooms and are not reported for the sake of space.

C. Experiment

Experiments have been performed to test the algorithms described in Sections III-A and III-B. Experiments have been conducted over periods of 10 hours from 8 pm to 6 am on different days in February and March 2015. During the experiments, the outside conditions were relatively consistent with outdoor temperature ranging from 4 to 10°C . In particular, since the outside temperatures were close to the ones experienced during the identification procedure, the identified steady-state model for outside temperature could be used. For the computation of the bid, it was assumed that the temperature at the beginning of the experiment is 23°C to allow a meaningful comparison between different days, and with the simulation results. Therefore, the temperature was regulated to this value before each experiment. Since the same model and initial condition was used in simulation and in the experiment, the result of the bidding problem were the same, as detailed in Section IV-A, with optimal bids that correspond to 25 and 43 % of the installed capacity, respectively.

For the experiments, different realizations of the AGC signal were used for testing. After the commitment of the bid and baseline, the computation of the control inputs, which determines how energy is split across the rooms, is performed with a time step of 15 minutes, therefore allowing to compensate partly for forecast errors. In practice, the frequency of update of the AGC signal is faster than 15 minutes, but the controller computed can be used to apply control actions at a faster rate. A rate of one minute was used in the experiments. A Kalman filter is used to estimate the state of the system.

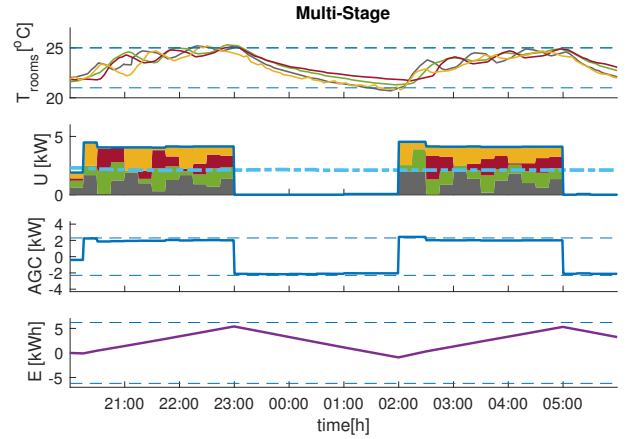


Fig. 7. Open-loop validation of battery multi-stage approach applying an "extreme signal extracted from the polytope Ξ_{ms} . Upper: Temperature variation for each zone. Middle Up: Power distribution among zones. Middle Down: Extreme signal variation and capacity bid. Lower: Extra energy stored.

After computing optimal bids solving respectively Problems (8) and (11), 4 closed-loop experiments were run, applying two different values of the AGC signal. Results are reported in Figure 6. For each experiment, four subplots are shown. The first one shows the evolution of the temperature in the four rooms, the second depicts the total power consumption in the four rooms and how it is split between the rooms. It can also be observed there how the energy dispatch in the four rooms is re-adjusted in closed loop every 15 minutes. The third plot shows the scaled AGC signal that needs to be tracked and the fourth plot shows the integral of the AGC signal over time, which represents the energy stored in the system as a result of the tracking. In the case of the two-stage method, the computed bid is higher and, therefore, results in larger tracking requirements which drive the temperature closer to the comfort limits. This confirms the results obtained in simulations. Small constraint violations are observed in the case of the two-stage method. This is expected since already in the case of perfect predictions and no model mismatch in simulations, the two-stage method displays an "aggressive" behaviour and runs very close to the constraints. The magnitude of those violations is however below 0.5°C .

It has already been highlighted that the conservatism of the multi-stage approach is partly due to the fact that it makes sure the system can track specific realizations of the AGC signal that are not very likely to be met in practice. One such "extreme" signal was chosen and another closed loop experiment was run. Results are shown in Figure 7. An interesting feature visible in this figure is the correlation between the temperatures in the rooms and the integral of the AGC signal (that can be thought of as the energy stored in the system as a result of tracking the AGC). It validates the idea that a battery is a reasonable proxy for the building in the sense that the building can also be "charged" and "discharged" with thermal energy over time to provide flexibility.

VI. DISCUSSION ABOUT THE EXPERIMENTAL SETUP

A. Applicability to HVAC systems

1) *Electric Heating*: The electric heaters were chosen for the following reasons: 1) Modelling electric heaters is simple

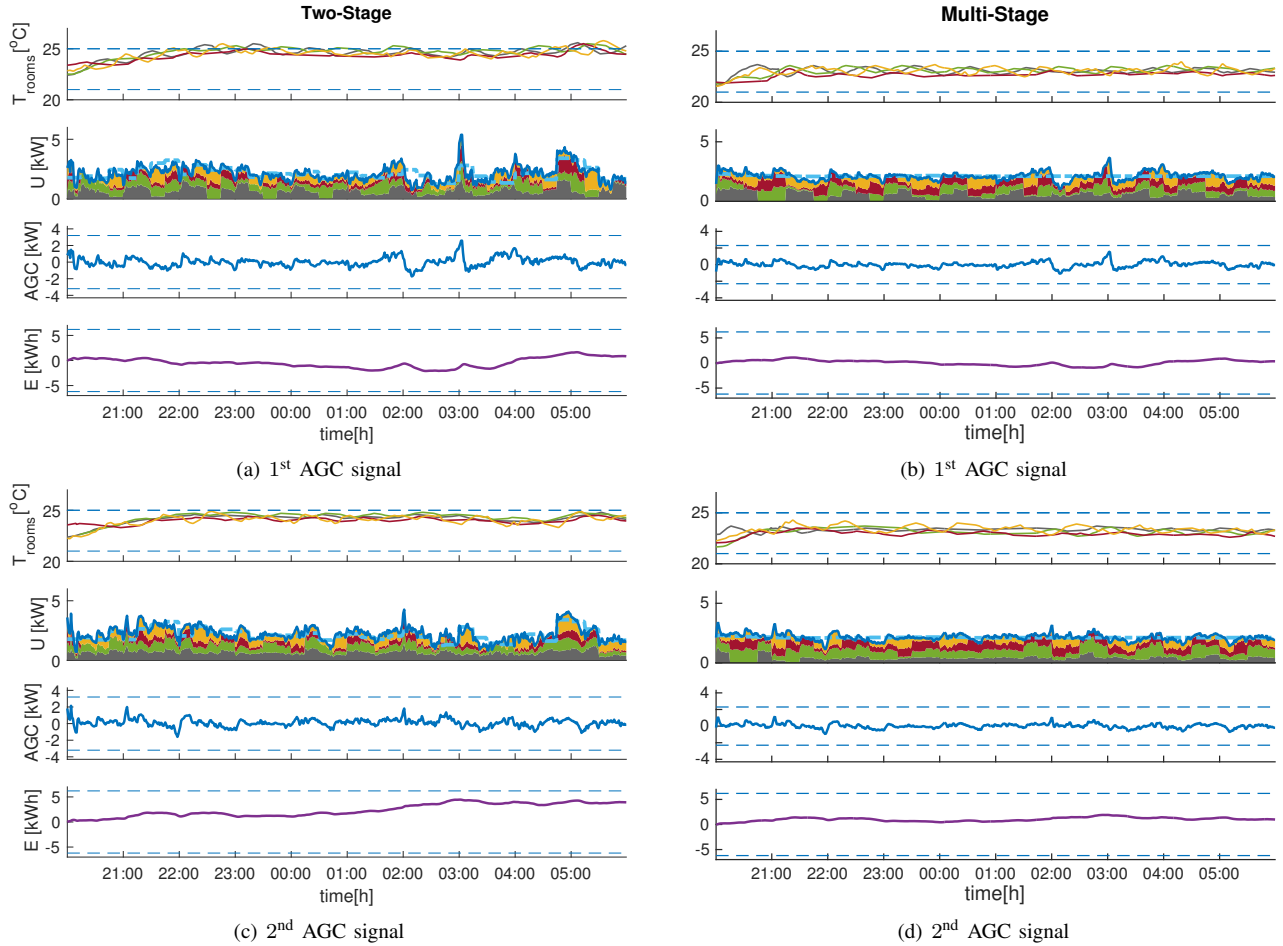


Fig. 6. Four experiments of AGC tracking. Two selected AGC signals, extracted from real data of 2013, are used to test and compare the two proposed controllers. Upper: Temperature variation for different zones. Each color corresponds to the measured temperature in each zone. Middle Up: Power distribution among zones. Middle Down: AGC signal variation and capacity bid. Lower: Integral of the AGC

2) It is possible to measure their power consumption directly, therefore, permitting to analyze the validity of the results. 3) They are highly responsive elements which permits to vary their power consumption at high frequency. 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 for residential and commercial buildings in Switzerland. In particular, recent federal statistics have revealed the presence of a quarter million electric-based heating units accounting for 4% of the total Swiss electricity consumption [23].

2) *Nonlinear HVAC dynamics*: In general, the main practical issue is to obtain a predictive model describing the conversion between electric power to the thermal state of the building as in V-B. This is a challenging task and represents one of the main limitations in the applications of MPC applied to building climate control. Nevertheless, many research efforts have been made to address this problem and there exist different modeling procedures based on first principle models. The reader is referred to [8] and references therein for more details. On the other hand, our experiments have shown that the relationship between thermal power input and zone temperatures can be identified with limited data successfully, and can be also supported by the use of software tools such

as detailed in [24]. While the model of the system does not in principle need to be linear, the methods rely on a linear model in order to form tractable reformulations of Problems (8) and (11). It is known that the response of typical HVAC systems may display nonlinear characteristics, especially on fast time scales. It is important to notice that, in the method we propose, the model of the HVAC is not taking into account the fastest dynamics. The model should be able to predict the behaviour of the HVAC on medium and slow time scales (5 to 15 minutes and above), which usually displays milder nonlinearities. Nevertheless, many methodologies, mainly based on sequential linearization or SQP approaches, have been proposed in the literature to tackle the complexity of the original non-linear, non-convex problem [9], [25], [26].

3) *Required Supporting Elements*: As it is in general not possible to modulate the power consumption of standard HVAC equipment at a fast pace, the last missing element is a resource capable of absorbing the fastest components of the tracking signal. In our case, the electric heaters are capable of that, but it has been shown in [7] that variable frequency drive fans can be used to provide fast modulation without affecting the rest of the HVAC system significantly. Another viable possibility is to couple the building with a storage element such as an electric battery. An investigation of the requirements

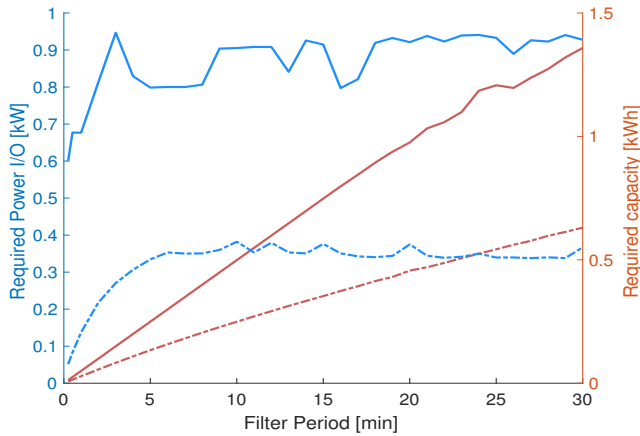


Fig. 8. Characteristics of the battery needed to absorb the fastest components of the AGC signal. Power ratings are in blue and energy capacity in red. Solid line is the worst case over one year, and dashed line represents the size needed to cover 99% of the signal. Values are computed for a $\pm 1\text{kW}$ offered tracking capacity.

for an electric battery needed to support the provision of the service in Switzerland is given in the following.

The yearly normalized AGC signal for 2014 is considered. First, we split the signal into a *fast* and a *slow* component. Second, we assume the building HVAC system to track the *slow* component of the signal, whereas the *fast* component is directly injected in the storage element. This process is repeated for several cutting frequencies mimicking different capabilities of the building HVAC. By doing this, we obtain requirements of the supporting battery both in terms of power and capacity specifications as a function of the cutting frequency. The results obtained using the battery model [27] are displayed in Figure 8.

VII. CONCLUSION

We presented an MPC framework for the provision of ancillary services using commercial buildings acting in the Swiss energy regulation scheme. We described how to formally assess the variation in power consumption that a building can support without impacting occupants' comfort. The effectiveness of the proposed approach has been demonstrated by means of simulation studies, and through experiments in a real multi-zone office building.

In this paper, we focused on overnight experiments. To extend the experiments to the day, more detailed modelling of the disturbances affecting the system, is needed, especially solar radiation which largely impact the system dynamics. Another research direction is to replicate the experimental results on a more typical and more complex HVAC system including for example chillers, heat pumps, fans.

ACKNOWLEDGMENT

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

REFERENCES

- [1] S. Meyn, M. Negrete-Pincetic, G. Wang, A. Kowli, and E. Shafieepour-fard, "The value of volatile resources in electricity markets," *Proceedings of the IEEE Conference on Decision and Control*, pp. 1029–1036, 2010.
- [2] B. J. Kirby, *Frequency Regulation Basics and Trends*, 2004, no. December.
- [3] N. Lu, "An evaluation of the HVAC load potential for providing load balancing service," *IEEE Transactions on Smart Grid*, vol. 3, no. 3, pp. 1263–1270, 2012.
- [4] H. Hao, T. Middelkoop, P. Barooah, and S. Meyn, "How demand response from commercial buildings will provide the regulation needs of the grid," *2012 50th Annual Allerton Conference on Communication, Control, and Computing, Allerton 2012*, pp. 1908–1913, 2012.
- [5] A. Costa, M. M. Keane, J. I. Torrens, and E. Corry, "Building operation and energy performance: Monitoring, analysis and optimisation toolkit," *Applied Energy*, vol. 101, no. 2, pp. 310–316, Dec. 2011.
- [6] M. Maasoumy, "Flexibility of Commercial Building HVAC Fan as Ancillary Service for Smart Grid," *IEEE Green Energy and Systems Conference (IGESC 2013)*, Nov. 2013.
- [7] Y. Lin, P. Barooah, S. Meyn, and T. Middelkoop, "Experimental Evaluation of Frequency Regulation from Commercial Building HVAC system," *IEEE Transaction on Smart Grid*, vol. 6, no. no. 2, pp. 776–783, March 2015.
- [8] D. Sturzenegger, D. Gyalistras, M. Morari, and R. S. Smith, "Model Predictive Climate Control of a Swiss Office Building : Implementation , Results , and Cost Benefit Analysis," *IEEE Transactions on Control Systems Technology*, vol. PP, no. 99, pp. 1–12, 2015.
- [9] F. Oldewurtel, A. Parisio, C. N. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and M. Morari, "Use of model predictive control and weather forecasts for energy efficient building climate control," *Energy and Buildings*, vol. 45, pp. 15–27, Feb. 2012.
- [10] H. Hao, A. Kowli, Y. Lin, P. Barooah, and S. Meyn, "Ancillary Service for the Grid via Control of Commercial Building HVAC Systems," *American Control Conference (ACC)*, pp. 467 – 472, 2013.
- [11] H. Hindi, D. Greene, and C. Laventall, "Coordinating regulation and demand response in electric power grids using multirate model predictive control," *Innovative Smart Grid Technologies (ISGT), 2011 IEEE PES*, no. January, pp. 1–8, 2011.
- [12] T. Gorecki, A. Bitlisliolu, G. Stathopoulos, and C. N. Jones, "Guaranteeing input tracking for constrained systems: Theory and application to demand response," in *the 2015 American Control Conference (ACC) (to appear)*, 2015. [Online]. Available: <http://arxiv.org/abs/1410.7297>
- [13] M. Maasoumy, C. Rosenberg, A. Sangiovanni-Vincentelli, and D. S. Callaway, "Model predictive control approach to online computation of demand-side flexibility of commercial buildings HVAC systems for Supply Following," in *Proceedings of the American Control Conference*, 2014, pp. 1082–1089.
- [14] M. Balandat, F. Oldewurtel, M. Chen, and C. Tomlin, "Contract design for frequency regulation by aggregations of commercial buildings," in *2014 52nd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, Sep. 2014, pp. 38–45.
- [15] E. Vrettos, F. Oldewurtel, and F. Zhu, "Robust Provision of Frequency Reserves by Office Building Aggregations," 2013.
- [16] I. Lympieropoulos, F. A. Qureshi, T. Nghiem, A. A. Khatir, and C. N. Jones, "Providing Ancillary Service with Commercial Buildings: The Swiss Perspective," in *9th IFAC International Symposium on Advanced Control of Chemical Processes (ADCHEM)*, Whistler, BC, Canada, June 2015.
- [17] Y. G. Rebours, S. Member, D. S. Kirschen, and M. Trotignon, "A Survey of Frequency and Voltage Control Ancillary Services Part I : Technical Features," *IEEE Transactions on Power Systems*, vol. 22, no. 1, pp. 350–357, 2007.
- [18] A. Shapiro, D. Dentcheva, and others, *Lectures on stochastic programming: modeling and theory*. SIAM, 2014, vol. 16.
- [19] Ben-Tal, A., El Ghaoui, L., and Nemirovski, A., *Robust Optimization*. Princeton University Press, 2009.
- [20] K. Margellos, P. Goulart, and J. Lygeros, "On the Road Between Robust Optimization and the Scenario Approach for Chance Constrained Optimization Problems," *IEEE Transactions on Automatic Control*, vol. 59, no. 8, pp. 2258–2263, Aug. 2014.
- [21] D. Kuhn, W. Wiesemann, and A. Georghiou, "Primal and dual linear decision rules in stochastic and robust optimization," *Mathematical Programming*, vol. 130, no. 1, pp. 177–209, 2011.
- [22] L. Ljung, *System identification*. Springer, 1998.

- [23] Office fédéral de l'énergie (OFEN), "Ex-post-analyse des schweizerischen energieverbrauchs 2000 bis 2013," Tech. Rep., sep 2014.
- [24] T. Gorecki, F. Qureshi, and C. Jones, "OpenBuild : An integrated simulation environment for building control," in *2015 IEEE Conference on Control Applications (CCA)*, Sep. 2015, pp. 1522–1527.
- [25] F. Oldewurtel, A. Parisio, C. Jones, M. Morari, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and K. Wirth, "Energy efficient building climate control using Stochastic Model Predictive Control and weather predictions," in *American Control Conference (ACC)*, 2010, Jul. 2010, pp. 5100–5105.
- [26] R. E. Griffith and R. Stewart, "A nonlinear programming technique for the optimization of continuous processing systems," *Management science*, vol. 7, no. 4, pp. 379–392, 1961.
- [27] O. Megel, J. L. Mathieu, and G. Andersson, "Maximizing the potential of energy storage to provide fast frequency control," in *IEEE PES ISGT Europe 2013*. IEEE, oct 2013, pp. 1–5. [Online]. Available: <http://ieeexplore.ieee.org/articleDetails.jsp?arnumber=6695380>



Luca Fabietti received his bachelor diploma in Information Engineering in 2011 and his master degree in Automation Engineering and Control in 2014 both from the University of Padova (UNIPD), Italy. In 2013 he joined the EES Smart Building group of Prof. Karl H. Johansson at KTH Royal Institute of Technology where he completed his master project. Since 2014 he is a Ph.D. student at the Laboratoire d'Automatique, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland. His research interests include optimal provision of high-

frequency demand and response control of virtual storage in Low Voltage and Medium Voltage grids with high penetration of uncertain renewable generation.



Tomasz Gorecki received his master degree in mechanical engineering with a minor in automatic control in 2012. He completed as master project a research project at the ONERA (French Aerospace public research Institute), studying application of model predictive control strategies for drones for area exploration. In September 2012 he joined the group of Colin Jones at the Automatic Control Laboratory at EPFL, Lausanne as a PhD student. He is employed in relation with the project entitled "Green energy management of structures" focusing

on application of Model Predictive Control to building and smart grid applications. His interests include robust and stochastic optimization, economic MPC, invariant set theory and computation, economics of the smart grid.



Faran A. Qureshi received his Bachelors degree in Electronic Engineering from the GIK Institute, Pakistan in 2009. He then received his Masters degree in Systems and Control from the Eindhoven University of Technology (TU/e), The Netherlands in 2012. In November 2012 he joined the group of Prof. Colin Jones in Automatic Control Laboratory at Ecole Polytechnique Federale de Lausanne (EPFL), Switzerland. He is working in the areas of model predictive control applied to smart buildings, ancillary service provision, and grid integration. His

research interests include economic MPC, stochastic, and robust optimization.




Altuğ Bitlislioglu received his diploma in Mechanical Engineering in 2011 from the Middle Eastern Technical University (METU), Ankara, Turkey. He then received his MSc. degree in Robotics, Systems and Control in 2013 from the Eidgenössische Technische Hochschule Zürich (ETHZ). In October 2013, he joined the group of Prof. Colin Jones in Laboratoire d'Automatique at École Polytechnique Fédérale de Lausanne (EPFL), as a PhD student. His research interests include assessment of tracking capabilities of distributed systems and optimal control methods

for smart grid applications.



Ioannis Lymperopoulos has received his bachelor and masters degree from the University of Patras in Greece at the department of Electrical and Computer engineering. He has obtained his doctorate from the Automatic Control Lab at the Swiss Federal Institute of Technology (ETH Zurich) and has continued as a post-doc at the Swiss Federal Institute of Technology in Lausanne (EPFL). His interests include data-based learning and control, constrained-optimization and non-linear filtering. He has been working on the fields of renewable energy and Ancillary Services

provision.



Colin N. Jones received the Bachelor's degree in Electrical Engineering and the Master's degree in Mathematics from the University of British Columbia, Vancouver, BC, Canada, and the Ph.D. degree from the University of Cambridge, Cambridge, U.K., in 2005. He is an Assistant Professor in the Automatic Control Laboratory at the École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland. He was a Senior Researcher at the Automatic Control Laboratory of the Swiss Federal Institute of Technology Zurich until 2010.

His current research interests are in the areas of high-speed predictive control and optimisation, as well as green energy generation, distribution and management.