

Multi-Time Scale Coordination of Complementary Resources for the Provision of Ancillary Services

Luca Fabietti^{a,*}, Faran A. Qureshi^a, Tomasz T. Gorecki^a, Christophe Salzmann^a, Colin N. Jones^a

^a*Automatic Control Laboratory, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland*

Abstract

This paper presents a predictive control scheme for coordinating a set of heterogeneous and complementary resources at different time scales for the provision of ancillary services. In particular, we combine building thermodynamics (slow), and energy storage systems (fast resources) to augment the flexibility that can be provided to the grid compared to the flexibility that any of these resources can provide individually. A multi-level control scheme based on data-based robust optimization methods is developed that enables heterogeneous resources at different time scales (slow and fast) to provide fast regulation services, especially a secondary frequency control (SFC) service. A data-based predictor is developed to forecast the future regulation signal, and is used to improve the performance of the controller in real-time operation. The proposed control method is used to conduct experiments, for nine consecutive days, demonstrating the provision of SFC service fully complying to the Swiss regulations, using a controllable building HVAC system on the EPFL campus and a grid connected energy storage system. The experimental results show that optimally combining such slow and fast resources can significantly augment the flexibility that can be provided to the grid. Moreover, by providing SFC service, the building can reduce its operational costs by up to 46% on average while maintaining a high level of occupant comfort. To the best of author's knowledge this work is the first experimental demonstration of coordinating heterogeneous demand-response to provide SFC service.

1. Introduction

In proportion to the total amount of electricity that is produced and consumed, very little energy storage capacity is available across the power grid. This is due to the very large costs of electrical energy storage devices which prevent their widespread deployment. As a consequence, production and consumption of electricity must practically be balanced at all times across any power network. However, due to inherently

^{*}This work has received support from 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 (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 755445).

*Corresponding author

unpredictable fluctuations, imbalances between the two are always present and need to be compensated in real-time to preserve the system frequency at its rated value (50 Hz in Europe) and, by doing so, prevent blackouts or the need to shed loads or generators from the network. This is typically achieved by keeping a set of power reserves on standby and ready to absorb any deviation. Depending on regional terminology, these reserves are called ancillary service providers (ASP), spinning reserves or operating reserves [1, 2]. These power reserves historically have been provided mainly by highly responsive power plants such as hydro-power plants, gas or coal stations, etc. However, utilizing these type of resources presents both economic and environmental issues [3]. Economic issues come from the fact that, in order to be able to respond to the network's need, ancillary reserves run at operating points that are not economically optimal, for example at part load. This incurs a loss-of-opportunity cost for ancillary reserves operators, for example a revenue deficit if a power plant runs at 90% of its full capacity, or conversely an excess cost when the power plant has to run while the market price of electricity is too low to cover the price of fuel. Environmental issues are due to the fact that fuel-based power plants are kept running as spinning reserves even when abundant renewable production is available. In addition, the increasing penetration of uncertain and uncontrollable sources of energy production such as solar and wind power have caused an increase in unpredictability and volatility of energy production.

It is therefore of paramount importance to integrate into the system new type of resources to improve the overall cost-efficiency of the power network. Potential candidates could be represented by, e.g., electrical storage systems (ESS) or thermal storage systems by controlling the heating, ventilation and air conditioning (HVAC) of large commercial buildings. ESSs are very well suited for ancillary services since they are highly controllable devices that exhibit very fast ramp rates [4]. In a landscape where the overall rotational inertia of the grid is decreasing, having such fast-responding Ancillary Services Providers (ASPs) could help reducing the frequency deviations and, thus, better stabilizing the operation of the electric grid. However, the main challenge when proving AS with ESSs is represented by the management of the State of Charge (SoC) level. In fact, the control signal to be tracked can exhibit significant biases over prolonged periods of time which can rapidly lead to the complete charge or discharge of the ESS. For this reason, in recent years different recharging strategies have appeared to optimize the provision of fast regulating services [3, 5, 6, 7, 8, 9]. All these studies typically focused on primary frequency control due to the smaller energy throughput that is typically required with respect to, e.g., Secondary Frequency Control (SFC). Nevertheless, due to the worst-case energy requirements and/or the conservative prequalification rules recently implemented by many Transmission System Operators (TSOs), even in this case, ESSs do not represent, in general, a economically viable solution due to very large capital costs [9].

On the other hand, the potential of demand-side resources to offer control reserves has been extensively recognized by both the academic and industrial community [10]. Buildings are responsible for 40% of the total energy consumption worldwide with roughly equal share for residential and commercial buildings [11].

Moreover, buildings are inherently characterized by a large thermal inertia that can be used to store large amounts of energy in the form of thermal energy. In particular, commercial buildings are good candidates for providing services to the grid for the following reasons: 1) they are typically characterized by a large HVAC system with respect to residential buildings. This corresponds to larger energy consumption per unit which in turn means a smaller cost of acquisition. 2) Most commercial buildings are already equipped with Building Energy Management Systems (BEMS) that can be readily used to monitor and control the operation of their HVAC systems [12].

However, the power consumption of a standard HVAC systems cannot be typically modulated at very high-frequencies due to physical limitations of the equipment. To mitigate this problem, several approaches have been proposed that typically focus on identifying specific HVAC components that can sustain such fast power changes. In [13], *Su et. al* propose a practical control framework to track a filtered version of the Area Generation Control (AGC) signal for secondary frequency regulation. This is achieved by acting on the chilled water supply setpoint of a chiller which, in turn, has a quantifiable effect on the electric power consumption of the HVAC system. A similar approach was also considered in [14, 15] where the power tracking was provided by adjusting the fan power consumption of the main air handling unit through either direct fan speed offset or by adjusting the mass flow setpoint. Also in this case, the building receives a filtered version of the AGC that is tracked in a $[1/30s \text{ to } 1/1min]$ or $[1/1min \text{ } 1/10min]$ frequency band, depending on the considered configuration. Despite displaying very fast response rates that are compatible with the typical SFC requirements, these methods have also some drawbacks. First, direct control of fan speed is not readily possible in most BEMS. Thus, this would require some level of retrofitting adding cost and complexity. Moreover, due to the complex control architecture of commercial buildings, slower control loops will likely compensate for net changes to supply pressure of mass flow which limits the ability of these strategies to track reference signals with slow time-scales. To the best of our knowledge, this is the reason why experiments of duration longer than 1h have not yet been demonstrated. A different approach has been proposed in [16, 17, 18], where authors propose to track the reference signal by adjusting the thermostat setpoint offset which has an indirect effect on the fan consumption through the corresponding change in the room damper and, therefore, the mass flow pressure. The advantages and disadvantages of this method are opposite of the previous ones. In general only software modifications would be required since thermostat changes can be done through many BEMS. However, since the electric power of fans is only controlled indirectly, communication and mechanical latency can significantly impact the tracking performance [18].

From the previous discussion, it is clear how ESSs and buildings are to some extent complementary resources. On one side ESSs are power-intensive devices with restrictive energy limitations; on the other side, buildings are energy-intensive devices with restrictive power limitations. Each type of resource could not provide ancillary services due to respectively economical and technical limitations but they have the potential to be operated together to provide reserves economically and reliably, and, in turn, improve

the overall efficiency of the network. Thus, the idea is to overcome the limitations of these resources by combining them into a single virtual resource. Thus, we would like the ESS to only take care of high-frequency components of the AGC while the building takes care of slower and more energy-intensive components.

Similar ideas have been recently explored more from a conceptual standpoint rather than technical. In particular, in [19, 20] the authors propose control frameworks to split the control signal at the TSO level so that the resulting components better fit the technical capabilities of different resources such as ESSs, supercapacitors, DR, etc. However, the implementation of these frameworks would require a drastic modification to the way the power system is currently operated. On the contrary, in the current paper, we propose to combine complementary resources to provide fast regulating services while complying with current regulations.

1.1. Contribution of the paper

In this paper, we provide a method to formally assess the aggregated flexibility that a set of heterogeneous resources such as an ESS and the HVAC system of a commercial building, can offer on the SFC market. Moreover, we present a multi-rate predictive control strategy to coordinate the two resources during real-time operation in order to collectively track the regulation signal received according to the contracted service.

The control methodologies are validated in a series of full-day experimental tests considering the aggregation of a fully-occupied office building on the EPFL campus and a grid connected ESS.

The contributions of the presented control framework with respect to existing works in the literature are:

1. By optimally splitting the AGC, the proposed method allows each resource to only track the components that best suit its technical capabilities. In particular, this allows virtually any type of building to participate in the provision of the service irrespective of the particular HVAC system. Conversely, by tapping the large thermal storage of the building, the method allows to drastically reduce the energy requirements on the ESS.
2. To better manage the status of the aggregated system, the closed-loop controller also adjusts the aggregated baseline power consumption on the intraday energy market.
3. The method guarantees at all times a high tracking quality without the need to filter the original signal. This is obtained by letting the ESS take care mostly of the high-frequency components of the AGC signal while the slower-frequency components are left to the building and the intraday market.
4. To further improve the closed-loop performance of the controller, a data-driven predictor for the future power requests has been developed that exploits the time-correlation and strong daily patterns of the AGC signal.
5. Experiments have been performed over extended periods of time (9 continuous days) which validates the robustness of the method against a wide range of uncertainty involved. Moreover, the experiments have been designed so as to be in complete compliance to the current regulation for the provision of

Name	Definition
EPFL	École Polytechnique Fédérale de Lausanne
ESS	Electrical Storage Systems
TSO	Transmission System Operator
AS	Ancillary Service
ASP	Ancillary Service Providers
AGC	Area Generation Control
SFC	Secondary Frequency Control
HVAC	Heating, Ventilation, Air conditioning
MPC	Model predictive controller
SoC	State of Charge
BEMS	Building Energy Management System
AGC	Area Generation Control
HLC	High Level Controller
LLC	Low Level Controller
HP	Heat Pump

Table 1: Acronyms.

SFC in the Swiss market. To the best of authors’ knowledge, this is the first work providing such an experimental validation over a prolonged period of time of the coordination of a set of heterogeneous resources proving SFC.

Notation: Throughout the article, \mathbb{R}^n denotes the n -dimensional real space, uppercase letters are used for matrices and lower case for vectors. a_k represents the value of vector a at time k , whereas bold letters are used to denote sequences over time, *e.g.*, $\mathbf{a} = [a_0^T, a_1^T, \dots, a_{N-1}^T]^T$. The notation \mathbf{P}^{res} indicates the real power flow of the particular resource, *res* whereas the bracket superscript notation, $\mathbf{P}^{\text{res},(j)}$ stands for the power trajectory corresponding the j -th scenario. Finally, the expected value operator over the probability distribution, ϵ , is denoted by $\mathbb{E}_\epsilon\{\cdot\}$. Please refer to Tables 1 for the list of all acronyms used in the manuscript.

2. Core Idea - Intuition

Fast regulating services to the electric grid, such as secondary frequency control (SFC), have been historically provided by traditional power plants such as hydro power plants, coal or gas stations, etc. The reason for this can be traced back to their relatively fast power responsiveness and to their inherent capabilities to absorb any energy bias of the regulation signal. The focus of this paper is to provide a control framework in order to reliably provide these kinds of services by combining ESS and commercial buildings. In the following section the core idea underpinning the whole proposition is provided.

2.1. Fast and Slow Resources

In this section we provide a qualitative description of the physical capabilities of both the ESS (fast) and the commercial buildings (slow resources) with respect to the typical requirements for the provision of SFC

to the grid.

In particular, we focus on a few key aspects that can be broadly categorized as *power requirements* and *energy requirements*. Following the framework proposed in [4], we compare the main characteristics of each resource with respect to the typical requirements for SFC (in the Swiss case). In particular, we focus on five key characteristics that are described in the following.

- Energy capacity: the maximum energy that can be shifted or stored by the resource.
- Ramp rate: the maximum rate at which the resource can modify its power production/consumption from its maximum to its minimum value or viceversa.
- Response granularity: the capability of the resource to implement any power setpoint between its minimum and maximum operating points.
- Response frequency: how often the power injection of the resource can be modified without wearing its physical equipment
- Response time: the maximum time elapsed between the power tracking request and the consequent modification of the power injection of the resource

To analyze the main characteristics of SFC with respect to the considered framework, we considered its frequency spectrum, shown in Figure 1. The spectrum is obtained from one year of historical realization of the AGC signal for 2014. A few observations are in order. First of all, high frequency components are quite damped due to the effect of both the system inertia as well as primary frequency control. Many distinct peaks are then visible in the medium range of the spectrum which corresponds to particular instants of the day and are mainly due to the way the market is operated. The highest peaks are at a frequency corresponding to 60, 30, and 15 minutes. More generally, an overall significant presence of low-frequency peaks can be observed, that are due to the integral action of SFC. Thus, a resource providing this service will need to track a signal spanning a wide range of frequency components (from 1/4 seconds to 1/60 min.) and with significant energy requirements.

Referring to Table 2, ESSs are highly responsive devices that exhibit reaction times and ramp rates that are only limited by the capabilities of their power inverters. Thus, ESSs not only meet, but even exceed, most of the requirements for the provision of SFC. However, due to the fact that the AGC signal can display significant biases over prolonged periods of time, the worst-case based dimensioning of ESSs typically represents a severe impediment to their widespread deployment. In fact, since the price of ESSs is typically determined by the required energy capacity and since the provision of SFC is rewarded in proportion to the power capacity offered, it is desirable to keep the energy/power ration as small as possible. Regarding commercial buildings, due to the large variety of different HVAC systems, it is difficult to provide a unique

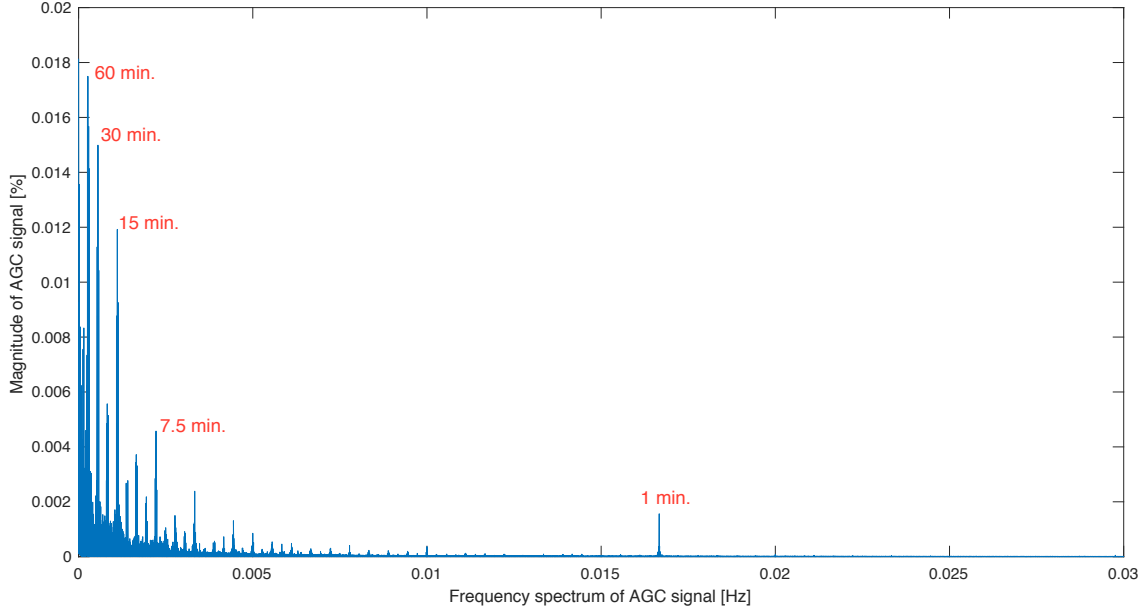


Figure 1: Spectrum of the AGC signal

identifier for all the considered key characteristics. Nevertheless, general qualitative statements can be given. First of all, commercial buildings are inherently characterized by a large thermal inertia that can be exploited to cheaply store energy without perceptibly affecting the occupants' comfort. On the contrary, HVAC systems are in general not suitable for fast regulating services [4]. In particular, their response time can be relatively slow due to the mechanical, control, and communication latencies that are introduced by complex BEMS [18]. Also strict constraints on ramp rates and response frequency are typically imposed to prevent an excessive stress and wear of the equipment.

Comparing the characteristics of these two resources, their complementarity is apparent. On one side, the ESSs are power-intensive devices with restrictive energy limitations; on the other side, buildings are energy-intensive devices with restrictive power limitations. Each of these resources cannot readily provide ancillary services due to respectively economical and technical limitations. However, they have the potential to be operated together to provide reserves economically and reliably, and in turn improve the overall efficiency of the network.

3. Resource Modeling

This section presents the mathematical models of all considered resources presented in the previous Section.

Property	SFC requirements	ESS	Commercial buildings
Energy capacity [kWh]	Significant	Limited (very expensive)	Significant (cheap)
Ramp rate [kW]	High	Very high	Moderate
Response granularity	Continuous	Continuous	Limited
Response frequency [Hz]	High	Arbitrary	Slow to fast
Response time [s]	High	High	Slow to Fast

Table 2: Qualitative description of the technical capabilities of ESSs and Commercial Buildings in relation to the requirements for the provision of SFC.

3.1. Building dynamics

We consider a commercial building served by an HVAC system that is either fully or partially controllable through a BEMS system.

3.1.1. Building Thermodynamics

The building thermodynamics are typically characterized by a set of linear equations of the form:

$$\begin{aligned} x_{i+1} &= f_b(x_i, u_i, d_i) \\ y_i &= f_z(x_i, u_i, d_i) \end{aligned} \tag{1}$$

where $x_i \in \mathbb{R}^{n_x}$ is the state of the system, $u_i \in \mathbb{R}^{n_u}$ is the thermal input to the building, $d_i \in \mathbb{R}^{n_d}$ is the disturbance input (outside temperature, solar radiation, occupancy, etc.), and $y_i \in \mathbb{R}^{n_y}$ is the temperature in different zones of the building at time step i . Typical sampling time is in the range of 15 minutes to 1 hour.

Thus, it is possible to define, over a specified horizon, N , a set of thermal trajectories that the building can support without violating its dynamics and constraints. More precisely, we define:

$$\mathcal{U}(x_0, \mathbf{d}) = \left\{ \mathbf{u} \left| \begin{array}{l} x_{i+1} = f_b(x_i, u_i, d_i) \\ y_i = f_z(x_i, u_i, d_i) \\ |y_i - T_{\text{ref}}| \leq \theta_i \\ u_i \in \mathbb{U} \\ \forall i = 0, \dots, N-1. \end{array} \right. \right\} \tag{2}$$

where x_0 is the initial state of the building, T_{ref} is a user-defined parameter which defines the ideal room zone temperature, and θ_i is the maximum allowed deviation, that can be time-varying, of the zone temperature from the ideal value. Thus, \mathbb{U} is the set defining the physical constraints on the thermal power input to the building which depends on the current weather prediction \mathbf{d} , and initial state.

3.1.2. HVAC System

The heating, ventilation, and air-conditioning system consumes the electric energy and produces the required thermal energy to maintain occupants' comfort in the building.

$$P_i^h = f_{hvac}(u_i, d_i) \quad (3)$$

We also define the set of feasible electrical power consumption profiles over the horizon length N :

$$\mathcal{P}^h(x_0, \mathbf{d}) = \left\{ \mathbf{P}^h \left| \begin{array}{l} P_i^h = f_{hvac}(u_i, d_i) \\ u_i \in \mathcal{U}(x_0, \mathbf{d}) \\ P_{min}^h \leq P_i^h \leq P_{max}^h \\ \forall i = 0, \dots, N-1. \\ T\mathbf{P}^h = \mathbf{0} \end{array} \right. \right\} \quad (4)$$

which is the set of electric power trajectories for the HVAC system for which both thermal and electric constraints over the prediction horizon are respected.

In the definition of \mathcal{P}^h , please notice the presence of a moving-block constraint (see e.g. [21]), $T\mathbf{P}^h = \mathbf{0}$. This constraint enforces that the electric power of the HVAC system is fixed over a certain number of time steps and is considered to encode the fact that, in general, it is not possible to adjust the power consumption too often either due to physical limitations of the equipment or to communication and mechanical latency that would inject significant fluctuations. Thus, the number of steps for which the power is blocked should be decided, depending on the particular application, in order to precisely track (on average) any given power setpoint. To provide an example, consider a building sampled at a 15 minute resolution for which the power consumption can be modulated only every 30 minutes. Assuming a horizon, N , equal to 6 steps, the matrix T would have the following form:

$$T = \begin{bmatrix} 1 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & -1 \end{bmatrix}$$

3.2. Electrical storage

We consider a grid-connected Electrical Storage System (ESS) for which the power injection can be modulated through a Battery Management System (BMS). We describe the temporal evolution of the SoC with the following difference equation.

$$\text{SOC}_{i+1} = f(\text{SOC}_i, P_i^s) \quad (5)$$

where SOC_i represents the SoC of the ESS at time step i , P_i^s its power injection, and $f(\cdot)$ is the map describing the time evolution of the SoC as a function of the current SoC and its power injection. In general, $f(\cdot)$ accounts for both conversion as well as temporal losses.

Both SoC and the power injection are constrained to lay within a feasible operational range at each time instance, i :

$$\text{SOC}_{\min} \leq \text{SOC}_i \leq \text{SOC}_{\max} \quad (6)$$

$$P_{\min}^s \leq P_i^s \leq P_{\max}^s \quad (7)$$

The set of feasible power of the ESS is defined as

$$\mathcal{P}^s(\text{SOC}_0) = \left\{ \mathbf{P}^s \left| \begin{array}{l} \text{SOC}_{i+1} = f(\text{SOC}_i, P_i^s) \\ \text{SOC}_{\min} \leq \text{SOC}_i \leq \text{SOC}_{\max} \\ P_{\min}^s \leq P_i^s \leq P_{\max}^s \\ \forall i = 0, \dots, N-1. \end{array} \right. \right\} \quad (8)$$

3.3. Total Power Consumption

As we are interested in combining these two resources into one single virtual entity, it is useful to introduce the total electrical power injection at time i , denoted as P_i , which is defined as the sum of the power injections at the two resources:

$$P_i = P_i^h + P_i^s \quad \forall i$$

Remark: Please note that we consider a passive sign notation, i.e., positive power values denote consumption.

4. Control Architecture

Referring to Figure 2, we provide herein an overview of the overall control architecture and how each component of it interacts with both the controlled resources as well as all external entities. The outputs and sampling times of the different modules in Figure 2 are given in Table 3.

Local Control & Measurements: We assume that each resource is already equipped with local sensors and a control system which monitors and controls the low-level functioning of the resource. In the case of BESS, this is typically represented by a BMS which provides SoC estimation, management of the individual cells as well as of local voltage and current constraints. Regarding commercial buildings, for most of them a BEMS is typically present which provides continuous monitoring and that can be used to control the power consumption of the HVAC by adjusting a number of user-defined setpoints.

More details for the particular experimental setup considered in this paper can be found in Section 6.

Forecasting Module: Responsible for the prediction of all uncertain quantities that affect the performance of the system. Thus, this module is in charge of obtaining both weather as well as AGC power requests predictions. These two quantities can be either locally computed, as in the case of the TSO future

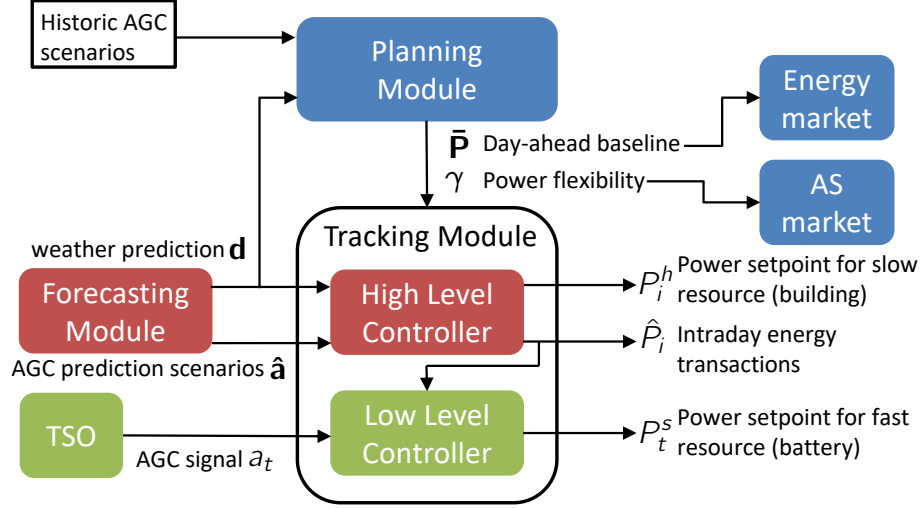


Figure 2: Overview of the control architecture. The modules in blue, red, and green colors operate with a sampling time of one day, 15 minutes, and 4 seconds, respectively.

power requests, or simply retrieved from external weather stations, as in the case of weather forecasts. The outputs of this module are sampled with a sampling time of 15 minutes.

More details can be found in Section 5

Planning Module: Activated once each day at a pre-specified time instant. Based on the most recent information, it solves an economic optimization problem to determine the aggregated power profile (day-ahead baseline) and the power capacity for the following day of operation. These quantities are then transmitted by means of an external secured channel to the Energy Market (baseline profile), and to the AS Market (power capacity) ¹.

More information can be found in Section 4.1.

Tracking Module: Activated during the delivery of the service. It is in charge of performing a multi-time scale coordination of all the resources so as to track the power request coming from the TSO while ensuring local constraints at all resources. The tracking module consists of two controllers - high level, and low level. The high level controller operates with a sampling time of 15 minutes and computes the setpoint for the slow resource (BEMS) and the intraday trades to modify the day-ahead baseline, while the low level controller operates with a sampling time of 4 seconds and computes the setpoint for the fast resource (BEMS).

Please refer to Section 4.2 for more information.

¹In this paper the power capacity is computed and transmitted daily. Clearly other possibilities can also be encompassed, e.g. weekly bidding

Module	Output	Sampling Time
Planning module	baseline power ($\bar{\mathbf{P}}$), and power capacity (γ)	1 day
Tracking Module		
High level controller	power setpoint for building (P_i^h), and intraday transaction (\hat{P}_i)	15 minutes
Low level controller	power setpoint for battery (P_i^s)	4 seconds
Forecasting module	AGC prediction scenarios ($\hat{\mathbf{a}}$), and weather prediction (\mathbf{d})	15 minutes

Table 3: Outputs and sampling times of the modules in Figure 2.

4.1. Planning Module

This section presents the Planning Module which is activated at a pre-specified time instant each day to plan the next day of operation. More precisely, the objective of this module is to determine the day-ahead baseline which is the power profile that the set of resources should track in the absence of regulation service (normal operation). The baseline power profile is denoted by the sequence \bar{P}_i , where the index i denotes a 15-minutes interval ² for the next day of operation. The module also determines a power flexibility, γ , which represents the maximum allowable deviation around the baseline that the aggregation of resources is willing to sustain for the next day of operation. This is done by solving the following economic optimization problem:

Problem 1 (Bidding problem).

$$\begin{aligned}
& \underset{\gamma, \bar{\mathbf{P}}, \pi_{P^h}, \pi_{\hat{P}}, \pi_{P^s}}{\text{minimize}} && \mathbb{E}_{\mathbf{a}} \left\{ J(\gamma, \bar{\mathbf{P}}, \hat{\mathbf{P}}, \mathbf{a}) \right\} \\
& \text{s.t.} && \\
& \text{(Building constraints)} && \mathbf{P}^h \in \mathcal{P}^h(x_0, \mathbf{d}) && \forall \mathbf{a} \in \Xi && (9) \\
& \text{(ESS constraints)} && \mathbf{P}^s \in \mathcal{P}^s(SOC_0) && \forall \mathbf{a} \in \Xi && (10) \\
& \text{(Total power)} && \mathbf{P} = \mathbf{P}^h + \mathbf{P}^s && \forall \mathbf{a} \in \Xi && (11) \\
& \text{(Power tracking)} && \|\mathbf{P} - \bar{\mathbf{P}} - \hat{\mathbf{P}} - \gamma \mathbf{a}\|_{\infty} \leq m_e && \forall \mathbf{a} \in \Xi && (12) \\
& \text{(Power flexibility)} && \gamma \geq 0 && && (13) \\
& \text{(Control Policies)} && \mathbf{P}^h = \pi_{P^h}(\mathbf{a}), \quad \hat{\mathbf{P}} = \pi_{\hat{P}}(\mathbf{a}), \quad \mathbf{P}^s = \pi_{P^s}(\mathbf{a}). && && (14)
\end{aligned}$$

where J represents the economic cost for the next day of operation which will vary depending on the specific AS regulations and country. In general, it would comprise different components such as the cost for the purchased baseline, the reward for the provided flexibility, a penalty cost for tracking violations, etc. A full overview of this cost function for the particular case of the Swiss market can be found in [22].

Equation (9) guarantees the satisfaction of comfort constraints as well as constraints on the HVAC system as detailed in Section 2.1. Similarly, (10) enforces the satisfaction of all constraints for the ESS. The

²The choice of the duration of this interval depends on the specific regulations for the country of interest. In the Swiss AS market the baseline must be specified for each 15 minute slot in the day.

total power consumption for the set of resources is defined in equation (11) whereas equation (12) imposes a minimum quality of tracking service. More precisely, it states that the aggregated power consumption during real-time operation, \mathbf{P} , should be adapted so as to guarantee a bounded tracking error:

$$|\epsilon_i| := |P_i - \bar{P}_i - \hat{P}_i - \gamma a_i| \leq m_e \gamma \quad i = 1, \dots, N$$

where \bar{P}_i is the committed baseline and the term \hat{P}_i is possible modifications of the committed baseline through intraday energy transactions (described in the following). Finally, the term $m_e \gamma$, determines the maximum allowable magnitude of tracking error in proportion to the offered capacity. This quantity is typically fixed as a participation requirement by the TSO.

The term \mathbf{a} , represents the AGC power requests, and it is, therefore, unknown at the time of decision. Thus, it is considered as an uncertain quantity belonging to the fixed uncertainty set Ξ which is constructed as detailed in following section. Due to the fact that the signal \mathbf{a} is revealed as time progresses, both the real-time power trajectories of the ESS and the building, as well as the intraday energy transactions are not determined once for all by the bidding problem. On the contrary, these quantities are encoded as causal functions of the observed uncertainty, i.e. the control action at each time instant, i , depends on all the past acquired information up to time $i - 1$. Finally, due to dependence on the uncertainty, the cost function, J , is evaluated in expectation.

Due to the infinite dimensionality of the decision space (policies), the problem as it is posed now cannot be solved. A scenario-based two-stage approximation scheme, similar to the one proposed in [23], is presented in Section 4.1.3 together with a specific form of the cost function J , and an explicit expression for the intraday policy, $\pi_{\hat{P}}(\mathbf{a})$.

4.1.1. Uncertainty set construction

To construct the uncertainty set, we consider, as is quite common, a scenario based approach. More precisely, the uncertainty set, Ξ , is constructed as the convex hull of a finite number N_s of past observed regulation signals and is given as

$$\Xi = \left\{ \sum_{j=1}^{N_s} \lambda^j \mathbf{a}^j \mid \sum_j \lambda^j = 1, \lambda^j \geq 0 \right\} \quad (15)$$

where \mathbf{a}^j are the previously observed AGC signal scenarios. [23] discusses the implications of the choice of Ξ in detail. The key idea is that, thanks to the statistical consistency of the AGC signal, if the controller is designed against past recorded instances of the AGC it should also be robust to future realizations.

4.1.2. Intraday Control

During real-time operation, the baseline power profile can be adjusted according to the current national market regulations. To provide an example, in Switzerland, it is possible to modify the baseline up to

30 minutes before delivery. When dealing with energy-constrained systems such as ESS or buildings, this is a very useful feature as it can be exploited to remove energy biases of the AGC signal that can have detrimental effects on the resources. Moreover, this can also be encoded into the bidding problem resulting in a larger amount of offerable service. To encode this into our bidding problem, we have considered an offline causal control policy as first introduced in [22]. For the sake of clarity, in the following, we simply report the policy.

For the combination of the resources offering the service, the residual tracking signal \mathbf{r} is the effective regulation signal with respect to the day-ahead declared baseline, after making the intraday adjustments. It is defined as the sum of the regulation signal and the normalized intraday control action

$$\mathbf{r} = \mathbf{a} + \hat{\mathbf{P}}^n$$

where $\hat{\mathbf{P}}^n$ is the normalized intraday control action, i.e., $\hat{\mathbf{P}} = \gamma \mathbf{P}^n$. The tracking constraint in (12) can be written in terms of the residual signal as

$$\|\mathbf{P} - \bar{\mathbf{P}} - \gamma \mathbf{r}\|_{\infty} \leq m_e$$

The idea of the intraday control policy is to modify the day-ahead baseline in such a way that it minimizes the energy content (cumulative sum) of the resulting residual tracking signal. A regulation signal with smaller energy content means that it will have a smaller bias, and will be closer to a zero-mean signal. It also means that it will require a combination of resources with smaller overall energy capacity to track the signal. In other words, a given set of resources can provide a higher flexibility when required to track a regulation signal with a small energy content. For these reasons, the intraday control policy is designed to reduce the energy content of the residual tracking signal by making appropriate intraday transactions at each time step. The intraday policy used in this study was developed in [22], and is presented here for the clarity of the manuscript.

$$\pi_{\hat{\mathbf{P}}^n}(\mathbf{a}) = \left\{ \hat{\mathbf{P}}^n \left| \begin{array}{l} \hat{P}_{i+1}^n = \\ \operatorname{argmin}_{\hat{P}_{i+1}^n} | r_0^i + \hat{P}_{i+1}^n + \mathbb{E}_a[\hat{a}_i^{i+1}] | \\ \text{s.t.} \quad r_0^i = r_0^{i-1} + \hat{P}_i^n + a_{i-1}^i \\ r_0^{-1} = 0, \hat{P}_0^n = 0, a_{-1}^0 = 0 \\ \forall i = 0, \dots, N-1. \end{array} \right. \right\} \quad (16)$$

where the cumulative sum of a signal \mathbf{r} from time step j to k is defined as $r_j^k = \sum_{i=j}^k r_i$. This causal intraday control policy is a function of the past received AGC signal and determines the future intraday trades. At each timestep, it measures the received AGC signal and optimizes the future intraday control action that minimizes the expected energy content of the future residual tracking signal. Note that the future AGC

signal is unknown at each time step, therefore an expectation is used in the cost function which can be evaluated using historic scenarios of the AGC signal.

The causal intraday control policy (16) can be used to compute a trajectory of intraday control actions corresponding to a given trajectory of AGC signal, and similarly, it can be used to generate scenarios of residual tracking signal corresponding to given scenarios of the AGC signal.

4.1.3. Bidding solution

The bidding problem can be approximated using the definition of the uncertainty set Ξ (15) by the following two-stage robust optimization problem:

Problem 2 (Tractable Bidding Problem Formulation).

$$\begin{aligned}
& \underset{\gamma, \bar{\mathbf{P}}}{\text{minimize}} && -\gamma \\
& \text{s.t.} && \\
& (\text{Building constraints}) && \mathbf{P}^{h,j} \in \mathcal{P}^h(x_0, \mathbf{d}) && \forall j = 0, \dots, N_s && (17) \\
& (\text{ESS constraints}) && \mathbf{P}^{s,j} \in \mathcal{P}^s(\text{SOC}_0) && \forall j = 1, \dots, N_s && (18) \\
& (\text{Total power}) && \mathbf{P}^j = \mathbf{P}^{h,j} + \mathbf{P}^{s,j} && \forall j = 0, \dots, N_s && (19) \\
& (\text{Power tracking}) && \|\mathbf{P}^j - \bar{\mathbf{P}} - \gamma \mathbf{r}^j\|_\infty \leq m_e && \forall j = 0, \dots, N_s && (20) \\
& (\text{Power flexibility}) && \gamma \geq 0 && && (21) \\
& && && && (22)
\end{aligned}$$

Please notice how the intraday control policy has been directly incorporated in the residual signal, \mathbf{r}^j as detailed in equation (16). Moreover, the control policies, $\pi_{P^h}(\mathbf{a})$ and $\pi_{P^s}(\mathbf{a})$ have been replaced by an implicit parametrization, as common for two-stage approximations [24], in which the optimization problem can select a different aggregated power trajectory (and different power trajectory for the ESS and the building) for each scenario \mathbf{r}^j . Finally, the general form of the cost function J introduced in Problem 1 has been simplified so that the problem only tries to maximize the offered power capacity. This was done mainly for the sake of simplicity. However, as detailed in [22], this is also desirable from an economic perspective as most of the economic benefits come from the reward for the offered capacity. We highlight that more complex cost functions can be easily incorporated.

The solution of the bidding problem results in the optimal value of the baseline power $\bar{\mathbf{P}}^*$, and flexibility γ^* .

4.2. Tracking Module

Every day at 00:00 the committed baseline and capacity, computed by the Planning Module, is activated and the set of resources are required to deliver the SFC service every 4 seconds. This is done by means of a Tracking Module that optimally coordinates the two resources so that their total power consumption is equal, within the allowed error bounds, to the sum of the total baseline and the AGC signal scaled by the committed power capacity.

The Tracking Module is composed of two controllers continuously working at different timescales, one at 15 minutes and the other one at 4 seconds³. On one side, the High Level Controller (HLC) decides the setpoints for the slow resource (building) as well as the future energy transactions to place in the intraday market. On the other side, the Low Level Controller (LLC) computes the power injection of the ESS so that the set of resources precisely track the received AGC signal. In the following sections the formulation of the two controllers is detailed.

4.2.1. High-level

The HLC operates at a 15 minutes time step. It is responsible of computing an adequate power setpoint for the building for the following 15 minutes so as to: 1) respect comfort and operational constraints of the building 2) guarantee that the ESS will be operated within its operational constraints. Moreover, depending on the current status of the resources and the current AGC predictions, the HLC can decide to sell or buy energy in the intraday market so as to either guarantee a high tracking quality and to improve the overall efficiency of the system. This is done by continuously running the following algorithm at every slow iteration k :

1. Retrieve the most recent weather forecast over the considered prediction horizon, N .
2. Form the predicted vector of disturbances, \mathbf{d} , over the prediction horizon.
3. Retrieve a set of possible scenarios, \mathbf{a}^j for the AGC signal from the Forecasting Module.
4. Estimate the current state of the building, x_k by means of a standard Kalman filter.
5. Retrieve the current state of the battery, SOC_k .
6. Solve the following MPC problem:

Problem 3 (Tracking MPC Problem).

$$\begin{aligned}
& \underset{P_0^h, \hat{P}}{\text{maximize}} && \mathbb{E}_{\hat{\mathbf{a}}} \{J_{comfort}\} \\
& \text{s.t.} && \\
& (\text{Building constraints}) && \mathbf{P}^{h,j} \in \mathcal{P}^h(x_0, \mathbf{d}) && \forall j = 0, \dots, N_s && (23) \\
& (\text{ESS constraints}) && \mathbf{P}^{s,j} \in \mathcal{P}^s(SOC_0) && \forall j = 1, \dots, N_s && (24) \\
& (\text{Total power}) && \mathbf{P}^j = \mathbf{P}^{h,j} + \mathbf{P}^{s,j} && \forall j = 0, \dots, N_s && (25) \\
& (\text{Power tracking}) && \|\mathbf{P}^j - \bar{\mathbf{P}}^* - \hat{\mathbf{P}} - \gamma^* \hat{\mathbf{a}}^j\|_{\infty} \leq m_e, && \forall j = 1, \dots, N_s && (26) \\
& (\text{State update}) && x_0 = x_k, SOC_0 = SOC_k && && (27) \\
& (\text{Power flexibility}) && \gamma \geq 0 && && (28) \\
& && && && (29)
\end{aligned}$$

7. Send the computed setpoint to the local controller of the building

³These sampling times are specific to the Swiss market and will, in general, depend on the specific regulations for the SFC provision, and for the intraday market

8. Wait for the next iteration

In Problem 3, the quantity $\hat{\mathbf{a}}^j$ represents the scenarios of the future AGC signal (over the prediction horizon) as generated by the AGC predictor (Section 5.2), the cost function $J_{comfort}$ is user-defined and it can comprise different metrics as, e.g., the comfort quality for the building occupants.

The solution of the MPC problem is the power setpoint, P_0^h , for the HVAC system which is sent to the local controller in charge of tracking this for the following 15 minutes. Moreover, the problem also determines the adjustment of the committed baseline profile, $\hat{\mathbf{P}}$, which are placed as energy trades in the intraday market. Note that the first few time steps of $\hat{\mathbf{P}}$ are fixed by the previous iteration of the MPC controller and not an optimization variable to make sure that the appropriate delay of the intraday market is respected (45 minutes in the Swiss market).

4.2.2. Low-level

The low-level tracking controller computes the control input for the fast resource (ESS). Every 4 seconds, it measures the actual power consumption of the slow resource which, in general, might have small deviations with respect to the setpoint as computed by the HLC. Based on this information, the LLC computes the power setpoint at each fast iteration, t , for the ESS as follows:

$$P_t^s = \bar{P}_t^* + \hat{P}_t^* + \gamma^* a_t - P_t^h \quad (30)$$

where \hat{P}_t^* is the intraday trade fixed by the HLC, a_t is the received AGC signal, and P_t^h is the measured power consumption of the slow resource (building) at fast time step t .

The computed power setpoint for the ESS is then transmitted to the BMS controller which will be responsible to implement this given that all physical constraints of the ESS are respected. If this is not the case, the power injection is rejected and a tracking error will occur.

5. Forecasting Module

5.1. Weather forecast

Weather forecasts can be typically obtained through available web-services by simply specifying the geographic location of interest. This requires one to select the closest weather station in the same bioclimatic zone as the installation. In the experimental demonstrator of Section 6, the weather station was selected from four available stations in a 3 Km radius, based on the historical quality of the data it provides. (source MeteoSwiss, meteostation *Lausanne freiland*, GPS coordinates $6^\circ 38.56'$ $46^\circ 33.33'$). The forecasting module retrieves the most updated forecasts every 15 minutes for the next 24h.

5.2. AGC Predictor

The ability to anticipate future values of the power requests dispatched by the system operator helps greatly in improving the average economic performance of the controller. We start by highlighting some characteristics of the AGC:

- The AGC signal follows a periodic pattern with recurring daily and hourly patterns, as already discussed on the basis of Figure 1.
- The empirical probability distribution of the AGC is non Gaussian, and time-varying
- The AGC signal is strongly correlated over time.

In the view of using forecasts of the AGC to solve the multi-stage real-time MPC Problem 3, the availability of a probabilistic forecast of the AGC is desirable, as opposed to a simpler point forecast. These probabilistic forecasts should capture the time-correlated nature of the AGC appropriately. In view of these elements, a method based on a variation of the non-parametric probabilistic forecasting method presented in [25] has been adopted and is presented in this section.

We denote with a_t the AGC power request at time t as the realization of a random variable A_t . Following the observation that the mean of the AGC presents a strong daily pattern, its distribution around that mean is consistent over time, we use the following assumption:

$$a_t = \bar{a}_t + \tilde{a}_t, \quad \forall t \quad (31)$$

where $\bar{a}_t \in \mathbb{R}$ captures the daily mean, so that $\bar{a}_t = \bar{a}_{t+k*24h}$ for all k and \tilde{a}_t originates from a single generating random variable denoted \mathcal{A} . The AGC is therefore generated by the sum of a daily mean and a single generating random variable \mathcal{A} .

We denote with f the probability density function of \mathcal{A} and F the corresponding cumulative distribution function. Assuming F to be strictly increasing, we define the quantile of \mathcal{A} at level α as $q^\alpha = F^{-1}(\alpha)$. A forecast of this quantile is denoted \hat{q}^α . A non-parametric forecast of the density function is formed by collecting quantile forecasts as:

$$\hat{f} = \{\hat{q}^{\alpha_i} | 0 \leq \alpha_1 < \dots < \alpha_m \leq 1\} \quad (32)$$

Based on m observed realizations $(\tilde{a}^{(j)})_{j=1,\dots,n}$ of \mathcal{A} , unbiased estimates of the quantiles are formed as the empirical quantiles of the observed realizations, ie:

$$\hat{q}^{\alpha_i} = \min_x \frac{\#\{\tilde{a}^{(j)} < x\}}{n} \geq \alpha_i \quad (33)$$

In turn, \hat{F} is obtained as the linear interpolation of the empirical cumulative distribution function between levels $(\alpha_i)_{i=1,\dots,m}$.

It is desirable that the estimated and observed quantiles are as close as possible, and should asymptotically match exactly. This property is referred to as reliability. When the estimates \hat{f} are reliable, the random variable Y whose realization at time t is defined as:

$$Y_t = \hat{F}(\tilde{a}_t) \quad (34)$$

is uniform, that is $Y \sim U([0, 1])$.

We will exploit the fact that predictions can be made Gaussian using a suitable transformation. We can obtain a normally distributed variable function X from Y as follows:

$$X_t = \Phi^{-1}(Y_t), \forall t \quad (35)$$

where Φ^{-1} is the probit function defined as $\Phi^{-1} : p \rightarrow \sqrt{2} \text{erf}^{-1}(2p - 1)$ and erf^{-1} the inverse error function. This yields $X_k \sim \mathcal{N}(0, 1)$, Gaussian-distributed with 0 mean and variance 1.

To capture the time correlation of the AGC, we assume for each time t that the random vector $\mathbf{X}_{t-K:t+N-1} = (X_{t-K}, \dots, X_{t+N-1})$ follows a multivariate Gaussian distribution $\mathcal{N}(\mathbf{0}, \Sigma_t)$, with K and N a horizon in the past and future that are long enough to capture the time-correlation of the AGC. It is assumed that N and K are fixed and we simply refer to $\mathbf{X}_{t-K:t+N-1}$ as \mathbf{X}_t . Once more, it is assumed based on observations that all \mathbf{X}_t 's originate from a single generating multi-variate Gaussian random variable \mathcal{X} . We can then estimate its covariance Σ by using past observations. An unbiased estimate is given by:

$$\hat{\Sigma} = \frac{1}{n-1} \sum_{t=1}^n \mathbf{X}_t \mathbf{X}_t^\top \quad (36)$$

The following procedure is then used to generate scenarios at time t .

1. The observations $\mathbf{a}_{t-K:t-1} = (a_{t-K}, \dots, a_{t-1})$ for the last K time steps are collected.
2. The corresponding realization of X_τ for $\tau = t-K, \dots, t-1$ are computed by removing the mean and transforming the result to normal as:

$$X_\tau = \Phi^{-1}(\hat{F}(a_\tau - \bar{a}_\tau)) \quad (37)$$

3. The parameters of the multi-variate Gaussian characterizing the prediction (X_t, \dots, X_{t+N-1}) : $\hat{\mu}_{t:t+N-1}$ and $\hat{\Sigma}$ are calculated by marginalizing the estimated distribution $\hat{\mathbf{X}}_t$ with respect to the observation vector $(X_{t-K}, \dots, X_{t-1})$, which gives:

$$\begin{aligned} \hat{\mu}_N &= \hat{\Sigma}_{NK} \hat{\Sigma}_{KK}^{-1} \mathbf{X}_{t-K:t-1} \\ \hat{\Sigma}_N &= \hat{\Sigma}_{NN} - \hat{\Sigma}_{NK} \hat{\Sigma}_{KK}^{-1} \hat{\Sigma}_{NK}^\top \end{aligned} \quad (38)$$

where $\hat{\Sigma} = \begin{pmatrix} \hat{\Sigma}_{KK} & \hat{\Sigma}_{NK}^\top \\ \hat{\Sigma}_{NK} & \hat{\Sigma}_{NN} \end{pmatrix}$

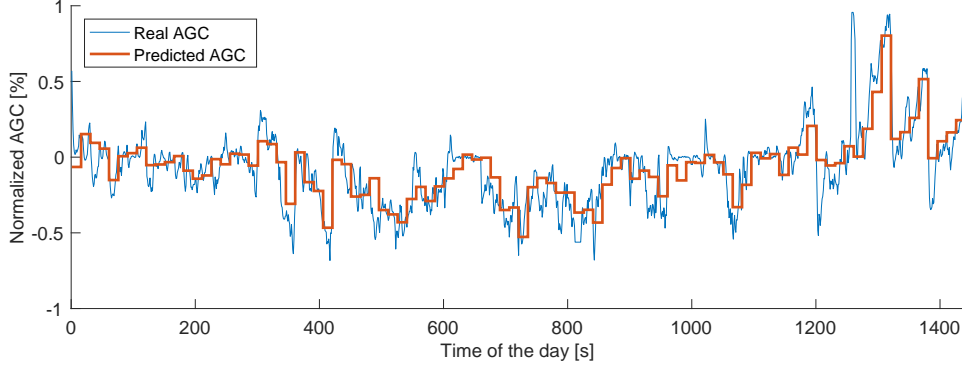


Figure 3: Experimental results on the EPFL campus: One full day of operation emulating the participation to the SFC in the Swiss Market.

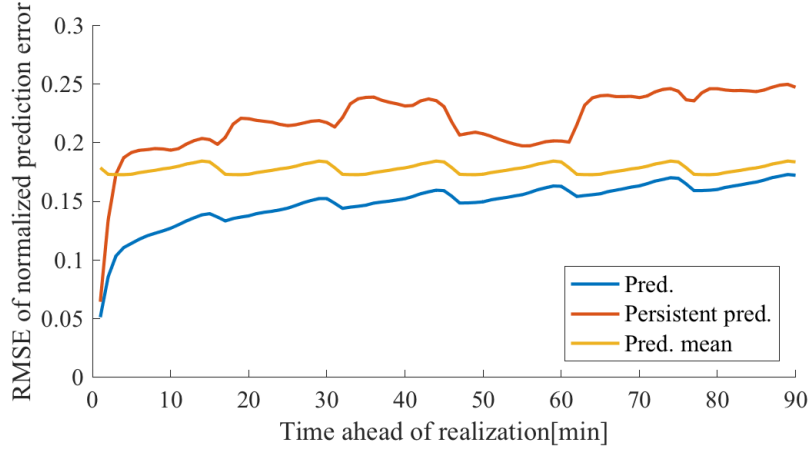


Figure 4: RMSE of the normalized prediction error as a function of lead time of prediction

4. N_s scenarios ($\mathbf{X}_{t:t+N-1}^{(i)} = (X_t^{(i)}, \dots, X_{t+N-1}^{(i)})$) $_{i=1, \dots, N_s}$ of this marginal are sampled in an iid fashion.
5. The inverse probit function is applied to each component of the scenarios to obtain $Y_\tau^{(i)} = \phi(X_\tau^{(i)})$, $\forall \tau = t, \dots, t+N-1$, $i = 1, \dots, N_s$.
6. The inverse of the estimated cumulative distribution function is finally used to compute the final forecast as $\hat{a}_\tau^{(i)} = \hat{F}^{-1}(Y_\tau^{(i)}) + \bar{a}_\tau$

This together provides forecasts for the AGC for the N next steps.

Figure 3 illustrates the performance of the predictor to predict the AGC. A one-day realization is depicted together with the 15 minutes-averaged prediction generated by the predictor and used in the MPC. We see that it can capture the trend quite successfully. Figure 3 shows the root mean square of the normalized prediction error comparing the predictor described in this section with two basic predictors. The mean Predictor predicts future values of the AGC to be the corresponding daily mean $\bar{\mathbf{a}}$. The persistent predictor predicts that the future value of the AGC is the current observed value, representing the fact the AGC is

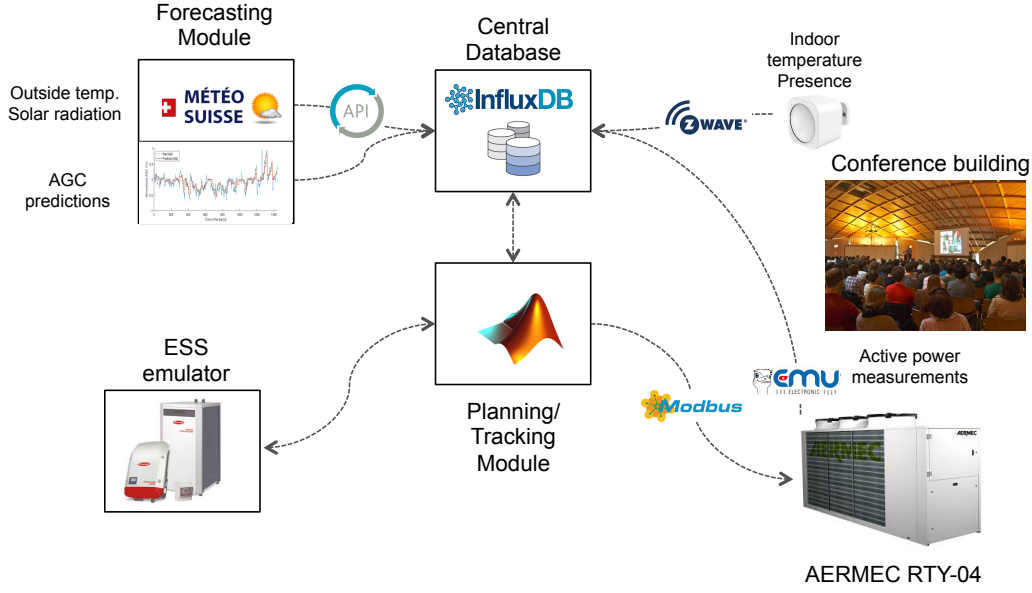


Figure 5: Overview of the experimental setup.

correlated over time. We see that our predictor achieves better predictions over the whole horizon, especially for horizons below 30 minutes.

6. Experimental setup

The experimental setup is depicted in Figure 5.

6.1. Commercial Building

We run our experiments in a relatively newly constructed building on the EPFL campus. The building of roughly 600 m^2 is used as a large audience/lecture room and is occupied on a regular basis with a maximum of 200 occupants. The building is served by a forced air-system that works for both cooling and heating depending on the season. More precisely, a single compressor roof-top heat pump, AERMEC RTY-04, is installed which accounts for 6kW active power at peak. A proprietary controller inside the heat pump continuously monitors the return air temperature coming from the building, it compares it with the reference temperature, and determines the operating point of all its active components (compressor, fan, etc.). We decided not to overwrite the logic of the controller since this has been specifically designed by the manufacturer to reduce the stress and wear on mechanical components of the heat pump. Moreover, in order to re-design this internal controller, a certain level of retrofitting would have been needed as not all relevant internal parameters are readily available.

Thus, the control input to the system is represented by the reference indoor temperature. This reference change will be tracked by the internal controller of the HP with a consequent effect on the electric power

consumption of the HVAC system. A rule-based control routine interfacing the HP controller to the Tracking Module was developed. The routine receives as input a power setpoint every 15 minutes from the solution of Problem (3) and returns a sequence (1 minute resolution) of indoor temperature setpoints that are sent to the internal controller of the HP so as to track the given power setpoint. This has the significant advantage of requiring, in general, only minor software modifications to the existing BEMS as most of them allow the remote control of thermostat setpoints.

For the measurements, the building has been retrofitted with three wireless Aeotec Multisensors that continuously monitor indoor temperature, and presence. Weather data were collected from a nearby weather station every 5 minutes as explained in Section 5, including measurements for outside temperature, horizontal global solar irradiance and weather observation. The monitoring and control of the heat pump is done through the serial communication protocol, Modbus. A central processing unit continuously receives all measurements and it uploads them into the database, InfluxDB, which is specifically designed to handle time-series data. The database is stored on the local network and it is connected to the open-source visualization platform, Grafana, which allows one to continuously supervise the overall functioning of the system. The Planning, Tracking, and Forecasting Modules are implemented in MATLAB, running on a computer on the local network that can access the latest measurements from the database and can send the temperature setpoints to the building and the power injection setpoints to the battery emulator.

6.1.1. Building model identification

The identification of the building model was performed using standard black-box linear identification techniques. In particular, we performed three weeks of open-loop experiments where, in order to excite the system dynamics, the temperature setpoint to the HP controller was modified using a mix of step and PRBS signals within a safe range of temperatures. Regarding the weather, a reasonably varying pattern was observed over the total duration of the experiments.

The experiments data set was divided into two chunks, one for identification (two weeks), and one for validation (one week). A linear sub-space identification approach was used [26]. A state space dimension of order $n_x = 3$ was found to be sufficient to appropriately describe the system dynamics. The model comprises three inputs, i.e., P^h , the electric power consumption of the HP, T_{out} , the outside temperature, and Q_{sun} the global horizontal irradiance. To simplify the notation, we introduce the external perturbation vector, $d := [T_{out}, Q_{sun}]$. The output, y , is the indoor temperature inside the building which is obtained as the average of all installed sensors. The identified model is of the following form

$$\begin{aligned} x_{k+1} &= Ax_k + B_u P_k^h + B_d d_k + K e_k \\ y_k &= Cx_k + e_k \end{aligned} \tag{39}$$

where the term e_k represents the noise component and K the disturbance matrix.

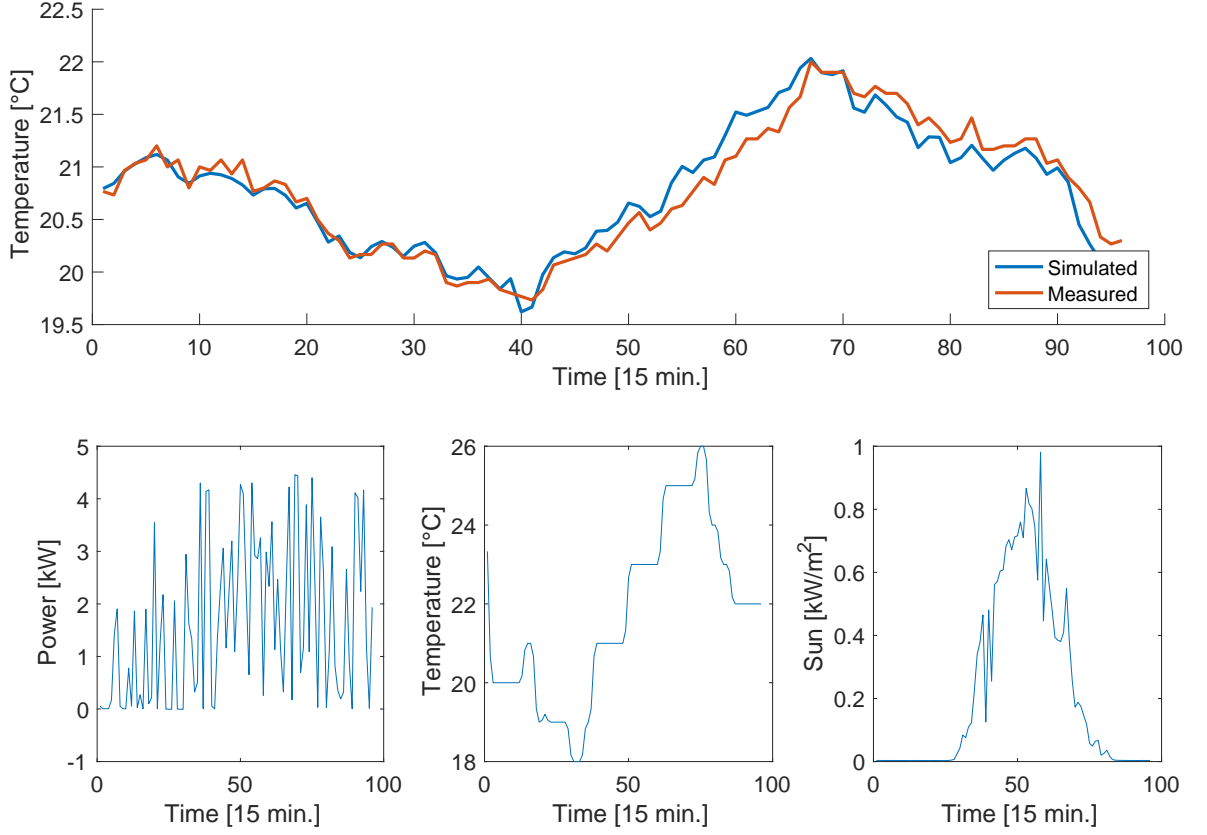


Figure 6: Validation of the model. The prediction error over one day is typically below 1 C. Upper: The actual measured indoor temperature (red) vs. the predicted temperature of the identified model. Lower left: The electrical power consumption of the HP. Lower central: The outside temperature. Lower right: Solar irradiance.

As can be observed in Figure 6, the model shows relatively good performance. For a quantitative evaluation, we consider the fit of the model which is computed as follows:

$$\text{FIT} := 100 \cdot \frac{1 - \|y - \hat{y}\|}{\|y - \bar{y}\|}$$

where \hat{y} is the predicted output, y the measured output, and \bar{y} the mean of the measured output.

The FIT ranges always between 70 and 83 % on a one day prediction horizon for all days in the validation dataset.

An inspection of the model reveals few key characteristics of the considered building. The impact of the solar irradiance is significant as the static gain from the solar input to the indoor temperature is $6.67^\circ\text{C}/(\text{kW}/\text{m}^2)$. The static gain from the outside temperature is less significant: $0.2^\circ\text{C}/^\circ\text{C}$. As a three dimensional state-space model was selected, three different time constants are present: one relatively fast at 15 minutes and two slower ones at 1 and 1.5 hours respectively. A simple physical interpretation can be

Quantity	Value
Energy capacity	5 kWh
Power limits	-5/5 kVA
η	0.95
a	1

Table 4: Technical specifications for the ESS considered in the experimental campaign.

given for this result: the fast time constant corresponds to the fast indoor air dynamics, whereas the slower dynamics can be associated to the walls and floor temperatures which constitutes for most of the thermal inertia of the building.

6.2. Electrical battery

In the experiments, a battery emulator was considered to propagate the battery SoC depending on the power injection as computed by the LL Controller (3).

In particular, at each fast sampling time (4 sec.), T_s the battery was simulated by means of the following non-linear model:

$$\text{SOC}_{t+1} = a\text{SOC}_t + \eta(\text{P}_t^s)_+ + \frac{1}{\eta}(\text{P}_t^s)_-$$

where $(\text{P}_t^s)_+$ represents the power injected in the ESS and $(\text{P}_t^s)_-$ power extracted from the ESS. Both the cycle efficiency coefficient, η , as well as the temporal losses coefficient, a , have been identified based on experimental data for the Leclanché large-scale ESS on the EPFL campus. More information regarding the modeling of the battery can be found in [27] whereas all relevant parameters implemented in the ESS emulator are reported in Table 4.

7. Experiments

In this section we describe the experimental campaign that was conducted in August 2017. To fully test the robustness and reliability of the controller, we run a set of nine days of continuous experiments emulating the participation of the combined system (BESS + building) into the Swiss SFC according to the current regulations. Thus, we consider historic AGC for the year 2013/2014 that was obtained from the Swiss TSO, Swissgrid. This large set of data is split into two subsets. The first subset is used to construct the uncertainty set (15) used in the Planning Module as described in Section 4.1.1. The second subset is used to randomly pick a realization of the AGC signal that is used during the closed-loop operation.

Everyday day at 23:00, the Planning Module collects the most current weather predictions from the nearest MeteoSuisse weather station. The weather forecasts comprise outside temperature, and solar radiation over a 24h period. The Planning Module further retrieves the current SoC level of the BESS and the most

updated state estimate of the building thermodynamics model from a continuously running Kalman filter. It then solves the bidding problem (2), and determines the power baseline at a 15-minutes resolution, and power flexibility for the next 24 hours starting at 00:00 of the following day.

At 00:00 everyday, the delivery begins and the Tracking Module is activated. It receives from the Planning Module the committed power profile, $\bar{\mathbf{P}}^*$ and capacity γ^* . As described in Section 4.2, at each slow iteration (15 minutes), it collects the current weather predictions over the prediction horizon together with the current state for each resource. The Tracking Module also obtains a set of possible AGC scenarios from the AGC Predictor. It then solves the HL Control problem (3) to determine both the power setpoint P^h for the controllable building as well as the amount of energy to trade on the intraday market \hat{P} . Finally, the LL Controller (30) computes the power injection setpoint for the ESS so as to track the received AGC signal.

Figure 7 displays one full day of operation. In particular, the topmost plot shows the pre-computed baseline power profile as computed by the Planning Module, the AGC request scaled by the bid capacity, and the power realization of the set of resources, i.e., the HP and the ESS, during real-time operation. As can be observed, the Tracking Module is able to optimally coordinate the two resources so as to perfectly track the received AGC signal. This is done while, at the same time, respecting the physical limitation at each resource. In particular, it can be observed in the third and fourth subplot of Figure 7 how both the SoC and the indoor temperature of the building are within their operational constraints at all times. In the second subplot, the allocation of the power, as performed by the HL Controller of the Tracking Module, is displayed. In particular, it can be observed how the HL Controller determines, at the 15 minutes time-scale, the power setpoint to the local controller of the HP so as to simultaneously preserve comfort inside the building and reset the SoC of the battery. This is particularly evident, for instance, between minutes 250 and 400 where the SoC is quite close to its upper limitation. At this time, the HL Controller decides to increase the power consumption of the HP with respect to the baseline profile in order to have a net discharging effect on the ESS and, therefore, reset its SoC to a safer value.

7.1. Multiplier effect

The statistics of the offered capacity, γ , over all experimental days are reported in Table 5. The comparison of the offered capacity between the case of a single ESS and the combination of the ESS and commercial building is performed. Please note that the offered capacity for the combined system (ESS + HP) depends on external varying factors affecting the system such as weather, occupancy pattern, etc. For this reason, the offered capacity displayed in Table 5 has some variability from one day to the next. On the contrary, the offered capacity with the ESS alone is computed as in Problem 1, where the total power (12) is obviously modified as $\mathbf{P} = \mathbf{P}^s$. Thus, the offerable capacity is only limited by the operational constraints of the battery and does not depend on other external conditions. This explains why the offered capacity is always the

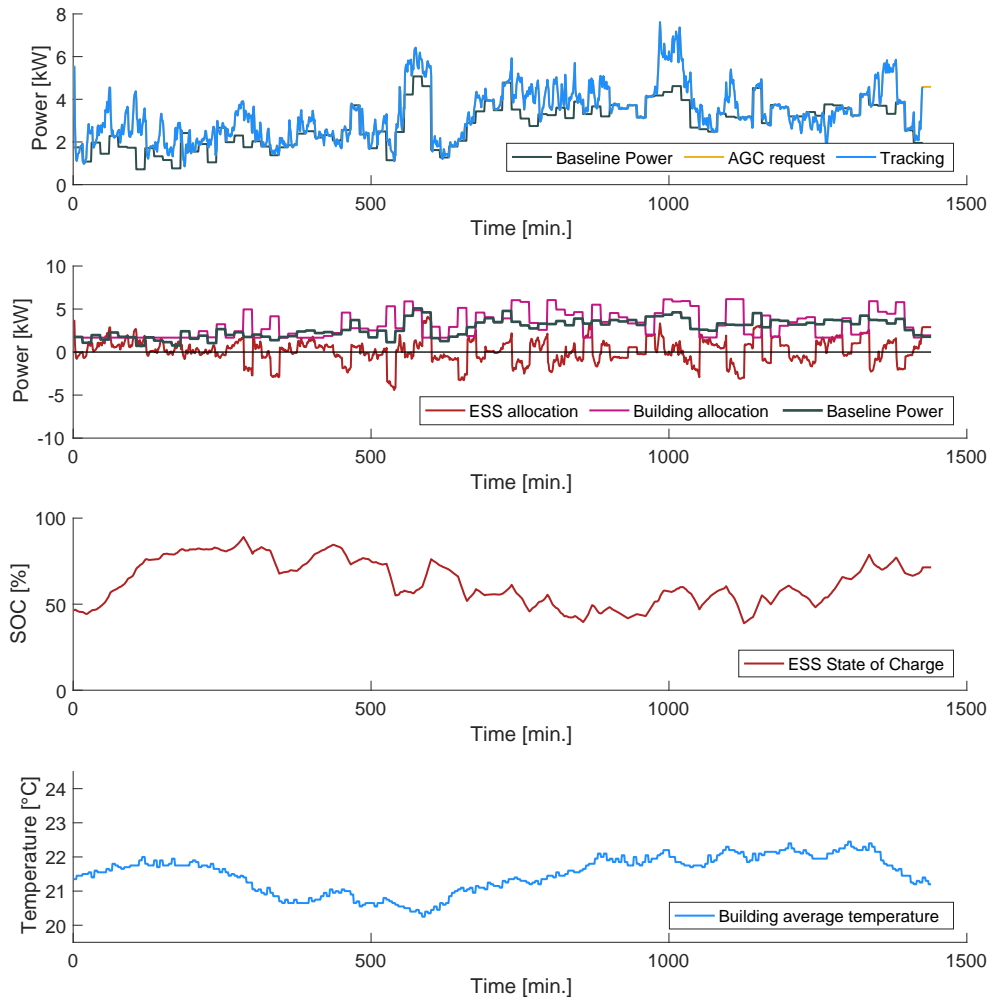


Figure 7: Experimental results on the EPFL campus: One full day of operation emulating the participation to the SFC in the Swiss Market.

Quantity	ESS	ESS + HP
Avg. capacity offered	3.4 [kW]	5.96 [kW]
Max. capacity offered	3.4 [kW]	6.72 [kW]
Min. capacity offered	3.4 [kW]	4.38 [kW]

Table 5: Statistics of the offered capacity over a period of 9 continuous days of experiments in August of 2017.

same each day. The results show that an overall substantial improvement is obtained through the proposed control method when combining a slow energy-intensive and a fast power-intensive resource in terms of the offered capacity to the grid. The experimental results confirm that exploiting the synergy between the slow and fast resources can increase the overall flexibility that can be provided to the grid.

7.2. Cost of Operation

We selected two sets of days for comparison, experimental days (**SFC experiments**), and the normal operation with no ancillary service provision (**normal operation**). To have a fair comparison between the two sets of data, we compare days for which weather conditions are very similar to the experimental days. In particular, we considered days for which the average outside temperature and solar radiation were in the same range as the minimum and maximum average values seen during the experiments. The building HVAC system during the normal operation days is operated using its default controller which is designed to stay within a small range around the ideal temperature (22 °C).

The building needs to interact with different energy markets while providing SFC. A detailed description of the energy markets and a comprehensive economic analysis for the provision of SFC was performed in our previous publication [22]. Therefore, the focus here is not to perform an in-depth economic analysis, but rather to get an estimate of the operating costs in both cases.

There are five components in the total operational cost when providing SFC service. (1) Baseline cost is the cost of buying the baseline energy in the day-ahead energy market. (2): Capacity Reward is the reward for providing the SFC capacity in the ancillary services market. (3): Tracking Penalty is the penalty for the errors in tracking the AGC signal. (4): Tracking bonus accounts for the fact that the extra energy consumed during tracking is at a reduced cost. (5): Intraday cost is the cost of modifying the day-ahead baseline by trading in the intraday energy market.

All the cost components for each experiment day are computed using the average Swiss price data for August 2014 to get an estimate of the total operating cost. The results are shown in Figure 8. For comparison, the total operating cost for the normal operation of the building using its default controller is also computed for several days, and the result is shown in Figure 9. The comparison shows that on average the operational cost is 5.94 CHF per day while providing SFC, and 11.08 CHF per day during normal operation. On average providing SFC results in about 46% reduction in the operating costs.

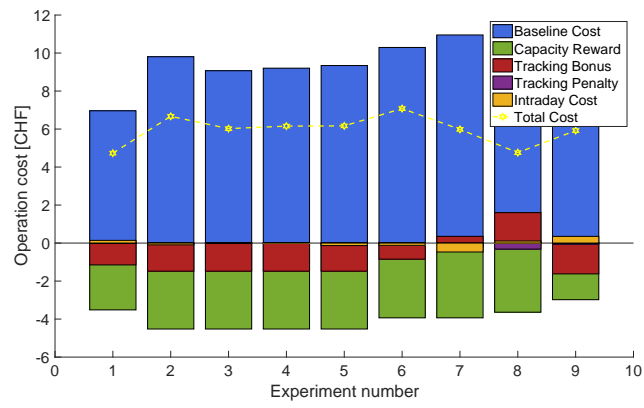


Figure 8: Break down of all economic costs for the 9 days of the experiment campaign.

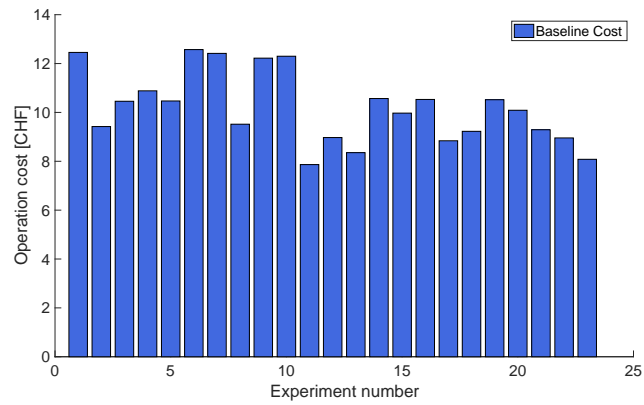


Figure 9: Baseline economic cost for the selected 23 days of normal operation.

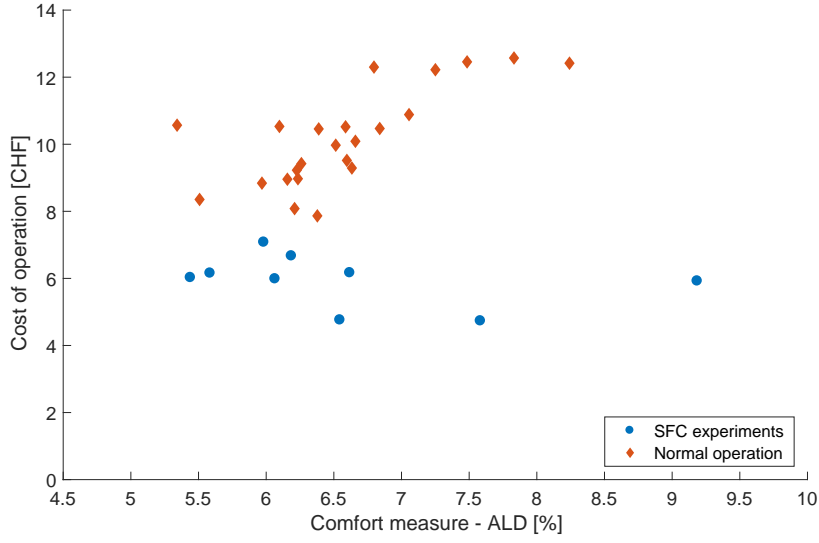


Figure 10: Cost of operation vs Comfort

7.3. Comfort

We studied the impact of providing SFC service to the grid on occupant comfort. To evaluate the occupant comfort, we used a standard measure, “Ashrae Likelihood of Dissatisfaction” (ALD) used in literature [28]. ALD is a function of the absolute difference between the room temperatures and an ideal (most comfortable) temperature defined by the occupants. It represents the average percentage of people dissatisfied in an environment. The lower the ALD value, the higher the occupant comfort.

The temperature trajectories of each day in the experimental sets (both SFC and normal operation) are used to compute the ALD comfort measure for each day. Moreover, the cost of operation is also computed, as explained in the previous section, taking into account all the cost components in accordance with the Swiss energy markets.

The results are depicted in Figure 10. Each day is denoted by a marker in the “cost of operation” against “comfort” plane. The blue markers represent the SFC experimental days, while red markers represent the normal operation days. The blue markers are always lower than the red ones. Given that each datapoint represents an entire day of operation, the total number of datapoints is limited and there is significant scatter in the data. Nevertheless, a clear pattern seems present and it can be concluded from the result that a similar level of occupant comfort can be provided at a reduced operating cost while providing SFC. In other words, the provision of SFC not only improves the operational cost of the building but also it does not affect the comfort level for the occupants.

8. Conclusion

This paper presented a framework for coordinating heterogeneous resources at multiple time scales to provide ancillary services to the grid. Particularly, ESS (fast) and building (slow resources) are considered in this paper, and are combined to utilize the synergy between the resources to maximize the flexibility provided to the grid. Mathematical models of the ESS, and the building were presented and a multi level control architecture was developed. The planning problem was formulated as a multi-stage uncertain optimization problem and was approximated using two-stage robust programming together with an offline causal intraday control policy. The tracking problem was divided into two levels, and an MPC controller was developed which at the higher level computed the power setpoint for the building, while a lower level tracking controller computed the setpoint for the ESS. The presented planning and tracking controllers were developed to handle the multi-time scale (fast and slow) nature of the resources. We also presented a data-based AGC predictor which was used in the control scheme.

The developed control method was validated on an experimental setup consisting of a fully-occupied controllable building on the EPFL campus, and an emulated grid connected ESS. The experimental results verify the efficacy of the proposed control scheme and shows that the building (slow resource) and ESS (fast resource) together can augment significantly the flexibility that can be provided to the grid than any of these resources can provide individually. Furthermore, an economic analysis shows that by providing SFC service, the building can reduce its cost of operation by up to 46% on average, without having any negative impact on the occupant comfort. To the best of author's knowledge this work is the first experimental demonstration of coordinating heterogeneous demand-response resources to provide secondary frequency control service.

Acknowledgment

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

References

- [1] B. Kirby, Ancillary services: Technical and commercial insights.
- [2] Y. G. Rebours, D. S. Kirschen, M. Trotignon, S. Rossignol, A survey of frequency and voltage control ancillary services mdash;part i: Technical features, *IEEE Transactions on Power Systems* 22 (1) (2007) 350–357. doi:10.1109/TPWRS.2006.888963.
- [3] M. Koller, T. Borsche, A. Ulbig, G. Andersson, Review of grid applications with the zurich 1mw battery energy storage system, *Electric Power Systems Research* 120 (2015) 128 – 135, smart Grids: World's Actual Implementations. doi:<https://doi.org/10.1016/j.epsr.2014.06.023>. URL <http://www.sciencedirect.com/science/article/pii/S0378779614002326>
- [4] F. Oldewurtel, T. Borsche, M. Bucher, P. Fortenbacher, M. G. V. T. Haring, T. Haring, J. L. Mathieu, O. Mège, 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, IEEE, 2013, pp. 1–24.
- [5] A. Oudalov, D. Chartouni, C. Ohler, Optimizing a battery energy storage system for primary frequency control, *IEEE Transactions on Power Systems* 22 (3) (2007) 1259–1266. doi:10.1109/TPWRS.2007.901459.

- [6] R. Hollinger, L. M. Diazgranados, C. Wittwer, B. Engel, Optimal provision of primary frequency control with battery systems by exploiting all degrees of freedom within regulation, *Energy Procedia* 99 (2016) 204 – 214, 10th International Renewable Energy Storage Conference, IRES 2016, 15-17 March 2016, Düsseldorf, Germany. doi:<https://doi.org/10.1016/j.egypro.2016.10.111>.
URL <http://www.sciencedirect.com/science/article/pii/S1876610216310724>
- [7] T. Borsche, A. Ulbig, M. Koller, G. Andersson, Power and energy capacity requirements of storages providing frequency control reserves, in: 2013 IEEE Power Energy Society General Meeting, 2013, pp. 1–5. doi:10.1109/PESMG.2013.6672843.
- [8] H. J. Kunisch, K. G. Kramer, H. Dominik, Battery energy storage another option for load-frequency-control and instantaneous reserve, *IEEE Transactions on Energy Conversion* EC-1 (3) (1986) 41–46. doi:10.1109/TEC.1986.4765732.
- [9] M. Koller, M. González Vayá, A. Chacko, T. Borsche, A. Ulbig, Primary control reserves provision with battery energy storage systems in the largest european ancillary services cooperation, in: Set of papers, CIGRE session 46 : 21-26 August 2016, Paris, CIGRE, 2016, pp. 361–NCA–C2–PS1, cIGRE Session 2016; Conference Location: Paris, France; Conference Date: August 21-26, 2016; .
- [10] H. C. Gils, Assessment of the theoretical demand response potential in europe, *Energy* 67 (2014) 1 – 18. doi:<https://doi.org/10.1016/j.energy.2014.02.019>.
URL <http://www.sciencedirect.com/science/article/pii/S0360544214001534>
- [11] A. Costa, M. M. Keane, J. I. Torrens, E. Corry, Building operation and energy performance: Monitoring, analysis and optimisation toolkit, *Applied Energy* 101 (2) (2011) 310–316. doi:10.1016/j.apenergy.2009.07.001.
- [12] M. Maasoumy, C. Rosenberg, A. Sangiovanni-Vincentelli, D. Callaway, Model predictive control approach to online computation of demand-side flexibility of commercial buildings HVAC systems for supply following, in: American Control Conference (ACC), 2014, 2014, pp. 1082–1089.
- [13] 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. arXiv:<https://doi.org/10.1080/23744731.2015.1072449>, doi:10.1080/23744731.2015.1072449.
URL <https://doi.org/10.1080/23744731.2015.1072449>
- [14] 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.
- [15] J. MacDonald, S. Kiliccote, Commercial buildings loads providing ancillary services in pjm, in: ACEEE Summer Study Energy Efficiency Buildings, 2014, pp. 1–5. doi:10.1109/PESMG.2013.6672843.
- [16] 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:<https://doi.org/10.1016/j.enbuild.2013.08.031>.
URL <http://www.sciencedirect.com/science/article/pii/S0378778813005288>
- [17] G. Goddard, J. Klose, S. Backhaus, Model development and identification for fast demand response in commercial hvac systems, *IEEE Transactions on Smart Grid* 5 (4) (2014) 2084–2092. doi:10.1109/TSG.2014.2312430.
- [18] I. Beil, I. Hiskens, S. Backhaus, Frequency regulation from commercial building hvac demand response, *Proceedings of the IEEE* 104 (4) (2016) 745–757. doi:10.1109/JPROC.2016.2520640.
- [19] P. Barooah, A. Buic, S. Meyn, Spectral decomposition of demand-side flexibility for reliable ancillary services in a smart grid, in: 2015 48th Hawaii International Conference on System Sciences, 2015, pp. 2700–2709. doi:10.1109/HICSS.2015.325.
- [20] T. Borsche, A. Ulbig, G. Andersson, A new frequency control reserve framework based on energy-constrained units, in: 2014 Power Systems Computation Conference, 2014, pp. 1–7. doi:10.1109/PSCC.2014.7038111.
- [21] R. Cagienard, P. Grieder, E. C. Kerrigan, M. Morari, Move blocking strategies in receding horizon control, in: 2004 43rd IEEE Conference on Decision and Control (CDC) (IEEE Cat. No.04CH37601), Vol. 2, 2004, pp. 2023–2028 Vol.2. doi:10.1109/CDC.2004.1430345.
- [22] F. A. Qureshi, I. Lympieropoulos, A. A. Khatir, C. N. Jones, Economic advantages of office buildings providing ancillary services with intraday participation, *IEEE Transactions on Smart Grid* PP (99) (2016) 1–1. doi:10.1109/TSG.2016.2632239.
- [23] L. Fabietti, T. T. Gorecki, F. A. Qureshi, A. Bitlislioglu, I. Lympieropoulos, C. N. Jones, Experimental implementation of frequency regulation services using commercial buildings, *IEEE Transactions on Smart Grid*.
- [24] A. Nemirovski, A. Juditsky, G. Lan, A. Shapiro, Robust stochastic approximation approach to stochastic programming, *SIAM Journal on Optimization* 19 (4) (2009) 1574–1609. arXiv:<https://doi.org/10.1137/070704277>, doi:10.1137/070704277.
URL <https://doi.org/10.1137/070704277>
- [25] P. Pinson, H. Madsen, H. A. Nielsen, G. Papaefthymiou, B. Klöckl, From probabilistic forecasts to statistical scenarios of short-term wind power production, *Wind energy* 12 (1) (2009) 51–62.
- [26] L. Ljung, *Ljung L System Identification Theory for User*, Vol. 25, 1987. doi:10.1016/0005-1098(89)90019-8.
- [27] F. Sossan, E. Namor, R. Cherkaoui, M. Paolone, Achieving the dispatchability of distribution feeders through prosumers data driven forecasting and model predictive control of electrochemical storage, *IEEE Transactions on Sustainable Energy* 7 (4) (2016) 1762–1777. doi:10.1109/TSTE.2016.2600103.
- [28] S. Carlucci, Thermal comfort assessment of buildings, Springer, 2013.