Title

Smart buildings in the smart grid: Contract-based design of an integrated energy management system

Authors

Maasoumy, M
Nuzzo, P
Sangiovanni-Vincentelli, A

Publication Date

2015

DOI

10.1007/978-3-662-45928-7_5

Peer reviewed

Smart Buildings in the Smart Grid: Contract-Based Design of an Integrated Energy Management System

Mehdi Maasoumy, Pierluigi Nuzzo and Alberto Sangiovanni-Vincentelli

Abstract In a supply-following "smart" grid scenario, buildings can exploit remotely controllable thermostats and "smart" meters to communicate with energy providers, trade energy in real-time and offer frequency regulation services, by leveraging the flexibility in the energy consumption of their heating, ventilation and air conditioning (HVAC) systems. The realization of such a scenario is, however, strongly dependent on our ability to radically re-think the way both the grid and the building control algorithms are designed. In this work, we regard the grid as an integrated, distributed, cyber-physical system, and propose a compositional framework for the deployment of an optimal supply-following strategy. We use the concept of assume-guarantee contracts to formalize the requirements of the grid and the building subsystem as well as their interface. At the building level, such formalization leads to the development of an optimal control mechanism to determine the HVAC energy flexibility while maximizing the monetary incentive for it. At the grid level, it allows formulating a model predictive control scheme to optimally control the ancillary service power flow from buildings, while integrating constraints such as ramping rates of ancillary service providers, maximum available ancillary power, and load forecast information. Simulation results illustrate the effectiveness of the proposed design methodology and the improvements brought by the proposed control strategy with respect to the state of the art.

Mehdi Maasoumy
Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720, e-mail: maasoumy@eecs.berkeley.edu

Pierluigi Nuzzo
Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720, e-mail: nuzzo@eecs.berkeley.edu

Alberto Sangiovanni-Vincentelli
Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720, e-mail: alberto@eecs.berkeley.edu

1 Introduction

Energy consumers do not usually pay enough attention to *when* they use energy. Demand for electricity tends to rise especially at times when it seems natural to use it; so natural, in fact, that we all tend to use energy at the same time and in similar ways. Such a common practice can easily lead to peaks in electricity demand that are traditionally met by operating extra power plants for limited portions of time, a solution which is generally expensive and environmentally unfriendly.

A more environmentally-friendly alternative to such a demand-following strategy is offered by *supply-following* (SF) programs, where utilities provide *incentives* to encourage consumers to reduce their demand during peak periods and use electricity at a less congested time. Another solution would be to significantly increase the penetration of Renewable Energy Sources (RESs), thus avoiding the introduction of expensive power plants. Several states in the United States and countries around the world have set ambitious targets for penetration of RESs by the next few years. The State of California, as an example, has targeted a 33% RES portfolio by 2020 [9]. However, a large-scale power grid requires continuous power balance between supply and demand; the power flows through individual transmission lines and facilities should also be controlled by continuously adjusting generation or load. Such an instantaneous matching becomes challenging due to the volatility, uncertainty, and intermittency of RESs, and makes their integration into the grid extremely difficult. The situation is even worsened by the uncertainties and randomness in the demand, due to short-term random switching of millions of individual loads, or longer-term (e.g. daily or seasonal) fluctuations in load and weather patterns.

While electricity storage is widely believed to be a solution to these problems by partially absorbing the variability associated with RESs and providing the extra power required at the peak energy demand hours, it is considered as an expensive and not environmentally-friendly solution. On the other hand, there is an emerging consensus that the transition from a demand-following strategy to a supply-following one will finally be enabled by the deployment of advanced metering infrastructure (AMI) devices in the *smart grids*. In a smart grid, new functions, denoted as *ancillary services*, are performed by the entities that generate, control, and transmit electricity in support of the basic services of generating capacity, energy supply, and power delivery. In this context, *smart buildings* can play a significant role. Buildings have inherent flexibility in the way their heating, ventilation and air conditioning (HVAC) systems consume electricity while respecting the occupants' comfort. This flexibility could be used to reduce costs if the electricity price is time-varying, or could be traded (i.e., sold to the utility) to be used for ancillary services.

Such flexible loads with thermal storage capabilities, also denoted as Thermostatically Controlled Loads (TCLs), are deemed to play an important role in regulating the grid frequency and, consequently, in enabling deep penetration of RESs. It has been reported that about 20% of the total electricity consumption in the United States is used by residential TCLs such as air conditioners, heat pumps, water heaters, and refrigerators [1, 2]. Recently, [20, 21] have shown that flexible loads such as TCLs are good candidates for providing ancillary services since their

aggregate flexibility can be controlled very fast, and sums up to tens of Gigawatts of power, only in the United States. Modeling, estimation, and control of aggregated heterogeneous TCLs for ancillary services have been discussed in [12]. TCLs are particularly well-suited for Direct Load Control and Demand Response programs that require loads to both decrease and increase power consumption because they are capable of storing thermal energy, much like a battery stores chemical energy. Despite several challenges in using loads for system services, key advantages include: (i) reduction in the overall grid emissions [32]; (ii) instantaneous response of loads to operator requests, versus slow response of generators to significant output changes [11]; and (iii) less variability associated to a very large number of small loads with respect to that of a small number of large generators [11]. Therefore, while ancillary services have been conventionally supplied by generation units, the increasing need for more energy storage capacity for frequency regulation, as well as more agile sources of ancillary services, makes it attractive to also use energy reserves on the demand side. It may soon be the case that the only technical impediment to reliable utilization of loads for system services is the development of the necessary models and control strategies and the development of inexpensive and scalable communication and sensing infrastructure [36].

To fully exploit the potential of buildings as service providers, we need to fundamentally re-design the way both the building HVAC system and the grid are controlled. This chapter addresses the problem of developing models and control algorithms for the deployment of a supply-following strategy in a smart grid, from the perspective of a hierarchical, distributed, cyber-physical system. The smart grid is intrinsically distributed, since different control algorithms must be executed in parallel on different components (e.g. buildings, energy providers) to achieve a common goal. On the other hand, the buildings can be abstracted as a load for the grid, which highlights the hierarchical nature of the system, where the designer is allowed to define the global behavior (via the grid controller) *together with* the local behavior of the "plant" (via the building controller).

Historically, very few, high-capacity reserves, such as industrial plants [33], have been used to provide ancillary services on the demand side. However, when a "swarm" of widespread and smaller capacity reserves are available, these service providers are better managed by intermediate entities called *aggregators*. The role of an aggregator is to provide appropriate incentives for a swarm of buildings at the right time, bundle the resulting capacity, and sell it in the wholesale market for frequency regulation. In this chapter, to simplify, we abstract into one *grid* agent all the players beyond the aggregator, such as the wholesale market players and the generation units, and denote as *buildings* the demand-side service providers that deal with the aggregator. We then focus on grid and buildings as the two sides of the supply-demand spectrum, by abstracting all the intermediate entities involved in the chain from power generation to power consumption.

To address the challenges originating from this distributed and hierarchical system, we resort to a Contract-Based Design (CBD) methodology. CBD has recently emerged as a compositional paradigm for the design of complex systems, emphasizing the concept of interface and requirement formalization to facilitate system

integration and provide formal support to the whole design flow [25, 31]. We use assume-guarantee contracts to formalize the requirements of the buildings, the grid, and their interface. Based on this formalization, we build on top of the supply-following scenario introduced in [20] and [21]. However, while in [20] and [21] the grid and building control schemes are derived and investigated separately, in this chapter, we provide an *integrated design framework* for Model Predictive Control (MPC) synthesis, which can combine and subsume both the approaches in [20] and [21]. At the building level, we develop an optimal control mechanism to determine the HVAC flexibility while maximizing the monetary incentive for it in a receding horizon fashion. At the grid level, we formulate a model predictive control scheme to optimally control the ancillary service power flow from buildings, while integrating constraints such as ramping rates of ancillary service providers, maximum available ancillary power, and load forecast. We use MPC as a convenient framework that allows optimizing a desired cost-function over a finite time horizon while, at the same time, satisfying a set of constraints.

The advantage of our contract-based methodology with respect to previous works is threefold: (i) it enables compositional design of the building and the grid MPC schemes, so that they can be independently implemented, while still guaranteeing that their integration is correct; (ii) it allows extending the approaches in [20] and [21] to highly distributed architectures, including a large number of control areas and buildings, in a scalable way; (iii) it supports automatic synthesis of embedded control software directly from assume-guarantee specifications.

The remainder of the chapter is organized as follows. Section 2 and Section 3 provide background information on the supply-following scenario of interest, and on our contract-based design methodology, thus setting the stage for our formulations. Section 4, Section 5 and Section 6 detail the main steps of our methodology, i.e. contract-based requirement formalization, generation of the model library, and MPC synthesis. Simulation results, in Section 7, illustrate the effectiveness of our design methodology and the improvements brought by the proposed control strategy with respect to the state of the art. Finally, Section 8 draws some conclusions.

2 A Supply-Following Scenario for Smart Buildings

In this section, we provide an overview of the Supply Following (SF) scenario considered in this chapter, by focusing on commercial buildings. Compared to residential buildings, commercial buildings typically have larger HVAC systems and therefore consume more electricity. In fact, commercial buildings account for more than 35% of electricity consumption in the US. Moreover, more than 30% of them have adopted a Building Energy Management System (BEMS) technology which facilitates the communication with the grid operators to provide flexibility. The majority of these buildings are also equipped with variable frequency drives, which in coordination with the BEMS, can modulate the HVAC system power consumption at intervals of the order of seconds. About 15% of electricity consumption in commer-

cial buildings is related to the fans of the HVAC systems. Fans are the main drivers, moving the conditioned air from the air handling units (AHU) to the rooms for climate control. For instance, the main supply fans that feed one of the buildings on the U.C. Berkeley campus, Sutardja Dai Hall, can spin at variable speeds, with the maximum rated power of 134 kW, proportional to the cube of the fan speed, which is about 14% of the maximum power consumed in the whole building. Moreover, the power consumed by the fans can be directly controlled upward or downward, thus making it an ideal candidate for ancillary services.

We refer to [15, 16, 18, 22, 28] for more information about the physics and control of HVAC systems. Moreover, we refer to [20, 21] for at-scale experiments on a real building, and a discussion on the feasibility of the proposed SF approach. We only observe here that modulating the fan speed of HVAC systems for extended periods of time with the existing control algorithms can lead to discomfort and does not allow optimizing the amount of flexibility provided by a building [19]. Hence, we propose to re-design the control algorithm, and consider commercial buildings whose HVAC systems are controlled by a Model Predictive Control (MPC) scheme running an optimal control problem at each time step k. Typically, the MPC aims at minimizing the total energy cost (in dollars). In an SF scenario, such cost must account for the reward received from the utility because of the building flexibility in energy consumption [20]. We refer to Table 1 and Table 2, later in the chapter, for a summary of the variables and parameters used in the description below.

To quantify the building *flexibility*, we adopt as a natural metric the difference between the upper and lower power envelopes that can be consumed without violating any constraints, i.e., at each time step k,

$$\text{Flexibility}(k) \triangleq P_f(e_k^u) - P_f(e_k^l) \tag{1}$$

where e_k^u and e_k^l are, respectively, the upper and lower bounds on the air mass flow of a building, and $P_f(.)$ returns the power consumption as a function of the air mass flow. Moreover, following the approach of [20], we assume that a *commercial contract* is stipulated between the utility and the building manager, whose duration, in terms of time steps, is $t_{ce} - t_{cs} = H^c$, where t_{cs} and t_{ce} are the commercial contract start and end times, respectively. The commercial contract has a limited duration because of the limited accuracy of the predicted flexibility by the building far ahead in time. Based on such commercial contract, the BEMS declares a lower envelope $\mathbf{e^l} = [e_{t_{cs}}^l, \ldots, e_{t_{ce}}^l]$, an upper envelope $\mathbf{e^u} = [e_{t_{cs}}^u, \ldots, e_{t_{ce}}^u]$ and a baseline $\mathbf{u}^* = [u_{t_{cs}}^*, \ldots, u_{t_{ce}}^*]$ air mass flow profiles for the duration of the contract. The utility is then allowed to select any power trajectory $P_f(\mathbf{u}) = [P_f(u_{t_{cs}}), \ldots, P_f(u_{t_{ce}})]$ such that, for all $k \in \{t_{cs}, \ldots, t_{ce}\}$, $P_f(e_k^l) \leq P_f(u_k) \leq P_f(e_k^u)$. However, the SF contract is deterministic, since the utility and the building operator both know how much money they have to pay or they receive from the beginning of the commercial contract. The utility charges the building operator for the baseline power consumption $P_f(\mathbf{u}^*)$, irrespective of the deviations due to flexibility signals, at a rate $\boldsymbol{\pi}^e$. On the other hand, the utility rewards the building operator for its declared flexibility, by

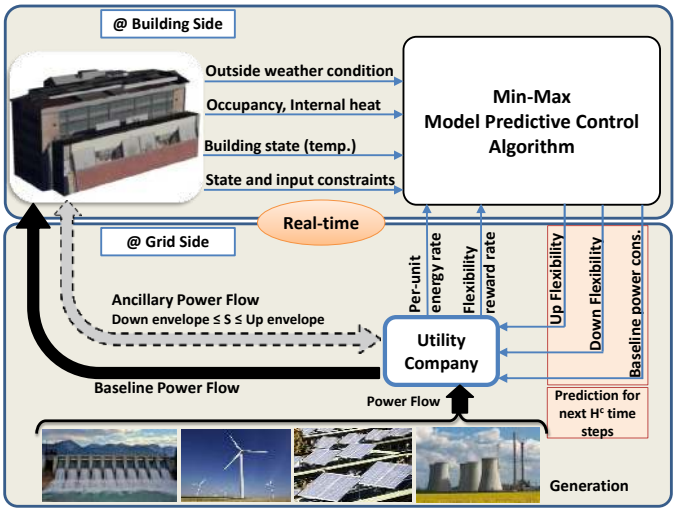

Fig. 1 Schematic of the proposed grid architecture and contractual framework.

providing both a downward flexibility rate $\underline{\boldsymbol{\beta}}$ and an upward flexibility rate $\overline{\boldsymbol{\beta}}$ with respect to the baseline power.

Clearly, by obeying the utility power consumption signals, the building may consume more or be in a worse state at the end of the H^c time slots with respect to a conventional demand-following protocol. The flexibility declared by the building operator would then be significant only if: (i) it is enough to be effectively exploited for frequency regulation services [21], and (ii) the reward from the utility is appropriate for the building. The schematic of the entire system architecture is shown in Fig. 1. The solid-line arrows correspond to the baseline power flow. The ancillary power flow is represented by a dashed-line arrow. For the architecture in Fig. 1 and the commercial contract summarized above, we can state our energy management integrated control problem as follows:

Integrated Energy Management Problem Statement. *Given* the real-time state of the buildings (e.g. indoor temperature, occupancy, internal heat, outside weather condition), a set of building temperature and control requirements, the real-time state of the grid (e.g. frequency deviation, generation and load forecast), a set of frequency regulation requirements, the per-unit energy price, the upward and downward flexibility rewards, and the duration of the commercial contract, *design* an optimal control strategy to determine the baseline power consumption, downward and upward building power envelopes, and grid flexibility signals for the buildings, while satisfying both the building and grid requirements.

In our scenario, the grid essentially controls the building consumption for the next H^c time slots, by sending flexibility signals (similar to frequency regulation signals) to be tracked by the HVAC fans as frequently as every few seconds.

In the next sections, after an overview of Contract-Based Design (CBD), we present a design framework that leverages a formalization of the control goals above and a library of models to generate the MPC schemes at both the building and the grid levels.

3 Contract-Based Design of Cyber-Physical Systems

The notion of formal contracts originates in the context of assume-guarantee reasoning. Informally, a contract is a pair $\mathcal{C} = (A, G)$ of properties, assumptions and guarantees, respectively representing the assumptions on the environment and the promises of the system under these assumptions. The essence of contracts is a compositional approach, where design and verification complexity is reduced by decomposing system-level tasks into more manageable subproblems at the component level, under a set of assumptions. System properties can then be inferred or proved based on component properties.

Compositional reasoning has been known for a long time, but it has mostly been used as a verification mean for the design of software. Rigorous contract theories have then been developed over the years, including assume-guarantee (A/G) contracts [4] and interface theories [7]. However, their concrete adoption in CPS design is still in its infancy [25]. Examples of application of A/G contracts have only been recently demonstrated in the automotive [5] and consumer electronics [26] domains. The use of A/G contracts for control design in combination with platform-based design (PBD) [30] was advocated in [25, 31], while in [27, 24, 10], a methodology was introduced that used contracts to integrate heterogeneous modeling and analysis frameworks for synthesis and optimization of CPS architectures and control protocols. The design flow was demonstrated on a real-life example of industrial interest, namely the design of system topology and supervisory control for aircraft electric power systems (EPS).

3.1 Contracts

We summarize the main concepts behind our methodology by presenting a simple contract model centered on the notion of platform *component*. A platform component $\mathcal{M}$ can be seen as an abstraction representing an element of a design, characterized by a set of *attributes*, including: *variables* (input, output and internal), configuration *parameters*, and *ports* (input, output and bidirectional); a *behavioral model*, uniquely determining the values of the output and internal variables given the values of the input variables and configuration parameters, and a set of *non-functional*

models, i.e. maps that allow computing non-functional attributes of a component, corresponding to particular valuations of its input variables and configuration parameters. Components can be connected together by sharing certain ports under constraints on the values of certain variables. In what follows, we use variables to denote both component variables and ports. A component may be associated with both implementations and contracts. An *implementation* M is an instantiation of a component $\mathcal{M}$ for a given set of configuration parameters. In the following, we also use M to denote the set of behaviors of an implementation, which assign a history of "values" to ports. Behaviors are generic and abstract. For instance, they could be continuous functions that result from solving differential equations, or sequences of values or events recognized by an automata model.

A *contract* $\mathcal{C}$ for a component $\mathcal{M}$ is a pair of assertions (A,G), called the *assumptions* and the *guarantees*, each representing a specific set of behaviors over the component variables [4]. An implementation M satisfies an assertion B whenever M and B are defined over the same set of variables and all the behaviors of M satisfy the assertion, i.e. when $M \subseteq B$. An implementation of a component satisfies a contract whenever it satisfies its guarantee, subject to the assumption. Formally, $M \cap A \subseteq G$, where M and $\mathcal{C}$ have the same variables. We denote such a *satisfaction* relation by writing $M \models \mathcal{C}$. An implementation E is a legal *environment* for $\mathcal{C}$, i.e. $E \models_E \mathcal{C}$, whenever $E \subseteq A$. Two contracts $\mathcal{C}$ and $\mathcal{C}'$ with identical variables, identical assumptions, and such that $G' \cup \neg A = G \cup \neg A$, where $\neg A$ is the complement of A, possess identical sets of environments and implementations. Such two contracts are then *equivalent*. In particular, any contract $\mathcal{C} = (A,G)$ is equivalent to a contract in *saturated form* (A,G'), obtained by taking $G' = G \cup \neg A$. Therefore, in what follows, we assume that all contracts are in saturated form. A contract is *consistent* when the set of implementations satisfying it is not empty, i.e. it is feasible to develop implementations for it. For contracts in saturated form, this amounts to verify that $G \neq \emptyset$. Let M be any implementation, i.e. $M \models \mathcal{C}$, then $\mathcal{C}$ is *compatible*, if there exists a legal environment E for M, i.e. if and only if $A \neq \emptyset$. The intent is that a component satisfying contract $\mathcal{C}$ can only be used in the context of a compatible environment.

Contracts associated to different components can be combined according to different rules. Similar to parallel composition of components, *parallel composition* ($\otimes$) of contracts can be used to construct composite contracts out of simpler ones. Let M_1 and M_2 two components that are composable to obtain $M_1 \times M_2$ and satisfy, respectively, contracts $\mathcal{C}_1$ and $\mathcal{C}_2$. Then, $M_1 \times M_2$ is a valid composition if M_1 and M_2 are *compatible*. This can be checked by first computing the contract composition $\mathcal{C}_{12} = \mathcal{C}_1 \otimes \mathcal{C}_2$ and then checking whether $\mathcal{C}_{12}$ is compatible. To compose multiple views of the same component that need to be satisfied simultaneously, the *conjunction* ($\wedge$) of contracts can also be defined so that if $M \models \mathcal{C}_1 \wedge \mathcal{C}_2$, then $M \models \mathcal{C}_1$ and $M \models \mathcal{C}_2$. Contract conjunction can be computed by defining a preorder on contracts, which formalizes a notion of *refinement*. We say that $\mathcal{C}$ refines $\mathcal{C}'$, written $\mathcal{C} \preceq \mathcal{C}'$ if and only if $A \supseteq A'$ and $G \subseteq G'$. Refinement amounts to relaxing assumptions and reinforcing guarantees, therefore strengthening the contract. Clearly, if $M \models \mathcal{C}$ and $\mathcal{C} \preceq \mathcal{C}'$, then $M \models \mathcal{C}'$. On the other hand, if $E \models_E \mathcal{C}'$, then $E \models_E \mathcal{C}$. Mathemati-

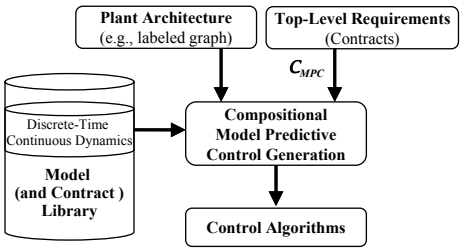

Fig. 2 Contract-based model predictive control synthesis flow.

cal expressions for computing contract composition and conjunction can be found in [4].

3.2 Design Flow

In [27, 24], we introduced a design methodology that addresses the complexity and heterogeneity of cyber-physical systems by using assume-guarantee contracts to formalize the design process and enable realization of control protocols in a hierarchical and compositional manner. Given the architecture of the physical plant to be controlled, the design is carried out as a sequence of refinement steps from an initial specification to a final implementation, including synthesis from requirements and mapping of higher-level functional and non-functional models into a set of candidate solutions built out of a library of components at the lower level. Initial top-level requirements are captured as contracts and expressed using linear temporal logic (LTL) [29] and signal temporal logic (STL) [23] formulas to enable requirement analysis and early detection of inconsistencies. Requirements are then refined into a controller architecture by combining reactive synthesis steps from LTL specifications with simulation-based design space exploration steps. We have demonstrated our approach on the design of embedded controllers for aircraft electric power distribution. In this work, we adapt and extend our methodology to the design of MPC algorithms. In our design flow, pictorially represented in Fig. 2, platform component design and characterization is completely orthogonalized from system specification and algorithm design.

Requirement Formalization. In the top-down phase of the design process, top-level system requirements are formalized as contracts. Responsibilities of achieving requirements are split into those to be established by a system (guarantees) and those characterizing admissible environments (assumptions). In a distributed control set-

ting as in our application, the requirements of a controller C can be expressed as a contract $\mathcal{C}_C = (A_C, G_C)$, where A_C encodes the allowable behaviors (i.e. trajectories, or sequences of valuations over a set of variables) of the environment (e.g. physical plant, or other controllers) and G_C encodes the required behaviors of C. To define $\mathcal{C}_C$, as shown Fig. 2, we can leverage a discrete time abstraction of the continuous behaviors of the components. We then express A_C and G_C as either first order difference equations involving the component variables and parameters (time varying properties), or arithmetic constraints on real variables that must hold at each time step (time invariant properties). The algebra of contracts can then be implemented by simply combining constraints via conjunction or disjunction to express, respectively, intersections or unions of behaviors. Examples of this approach will be provided in Section 4.

Platform and Contract Library Generation. In the bottom-up phase of the design process, a library of components (and contracts) is generated to model (or specify) both the plant architecture (e.g. the power system or the building) and the controllers. Components can be hierarchically organized to represent the system at different levels of abstraction, e.g. *steady-state* (static), *discrete-event* (DE), and *hybrid* levels. At each level of abstraction, components are also capable of exposing multiple, complementary *views*, associated with different design concerns (e.g. safety, performance, reliability) and with models that can be expressed via different formalisms (e.g. graphs, linear temporal logic, differential equations), and analyzed by different tools. Such models include non-functional (performance) metrics, such as timing, energy and cost. As detailed in Section 5, in this work, we model our platform components by adopting the same discrete-time abstractions and formalisms we use for requirements. System behavioral models are expressed using difference equations, while performance and cost models are polynomial functions of the component variables and parameters.

Mapping Functions to Implementations. System design (synthesis) is cast as a set of problems mapping functionality (specifications or requirements) over implementations. A mapping problem can be solved by casting an optimization problem that uses information from both the system and the component levels to evaluate global tradeoffs among components or minimize a cost function.

For an MPC scenario, let $\mathcal{C}_{MPC}$ be the contract formalizing the requirements of the closed loop architecture in Fig. 3, where M is the controller and P the plant, and let H be the MPC horizon. We assume that the system dynamics are described by the following difference equation at each time step t:

$$x_{t+1} = p(x_t, u_t, d_t) \quad \forall\, t \in \mathbb{N}, \tag{2}$$

where x_t is the system state, u_t the control input, and d_t an external uncontrolled input (disturbance). We denote as $s = (x, u, d)$ the set of system variables. A system behavior or trajectory $\sigma = s_0, s_1, s_2, \ldots$ is a sequence of valuations over s, for all $t \in \mathbb{N}$. We also assume that the assumptions and guarantees in $\mathcal{C}_{MPC}$ can be represented as follows:

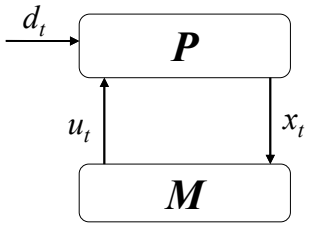

Fig. 3 Generic feedback control scheme.

$$A_{MPC} = \{\sigma | \alpha_t(s_t) \leq 0 \ \ \forall\, t \in \mathbb{N}\} \qquad G_{MPC} = \{\sigma | \gamma_t(s_t) \leq 0 \ \ \forall\, t \in \mathbb{N}\}, \tag{3}$$

where $\alpha_t(.)$, $\gamma_t(.)$ are generic real functions (constraints) parameterized by t. Finally, let $C(u,x)$ be a real function providing the system cost in terms of system state and control. Then the optimal control problem aiming at minimizing the cost over time horizon H while satisfying the system dynamics and contract can be formulated as follows:

$$\min_{\mathbf{u}_t} \quad \sum_{k=0}^{H-1} C(u_{t+k}, x_{t+k}) \tag{4a}$$

$$\text{subject to:} \quad x_{t+k+1} = p(x_{t+k}, u_{t+k}, d_{t+k}), \quad \forall\, k \in \{0,...,H-1\} \tag{4b}$$

$$\gamma_{t+k}(x_{t+k}, u_{t+k}, d_{t+k}) \leq 0, \quad \forall\, k \in \{0,...,H-1\} \tag{4c}$$

$$\alpha_{t+k}(x_{t+k}, u_{t+k}, d_{t+k}) \leq 0, \quad \forall\, k \in \{0,...,H-1\} \tag{4d}$$

where $\mathbf{u}_t = (u_t, u_{t+1}, \ldots, u_{t+H-1})$. Both contract assumptions and guarantees are captured as optimization constraints. The resulting optimal control algorithm executes the optimization problem (4) in a receding horizon fashion, and is returned as the final design.

4 System Requirement Formalization

We consider a control area network as the one shown in Fig. 4. In particular, for the sake of illustration, we use a simplified two-area model, and abstract the complexity of the full power transmission and distribution subsystem. Commercial buildings have HVAC systems controlled by an MPC scheme, while the grid utilizes a hierarchical control scheme composed of a high-level MPC framework on top of the low-level classical automatic generation control (AGC), detailed in Section 5.2. The MPC schemes are based on discrete-time models of the system dynamics. Let τ_G and τ_B be the sampling times for the grid and building dynamics, respectively.

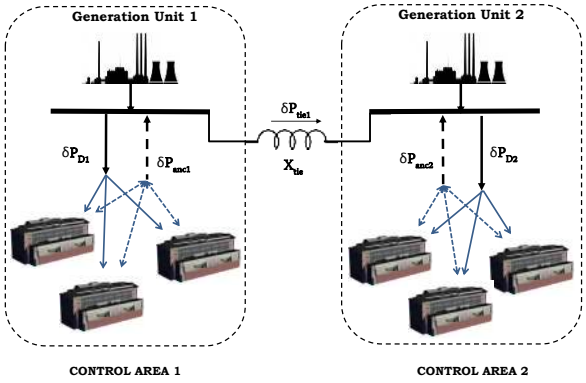

Fig. 4 A two-area model of the system architecture considered in this chapter.

Generally, τ_G is much smaller than τ_B since the grid must be able to send control signals as fast as every second, while τ_B may range from 15 min to 1 h.

We also assume that H^m and H are the time horizons of the MPC scheme adopted, respectively by the building subsystem (B-MPC) and the grid (G-MPC), with $H^m \gg H^c \gg H$, where H^c is the length of the commercial contract. The choice of H^m depends on how far in the future the predicted values of the building model parameters have an acceptable accuracy and the inputs to B-MPC (e.g. cost of energy) are available. Typical values for H^m range from a few hours to a few days. For instance, by assuming $\tau_B = 1$ h for the building dynamics, H^m may range from 3 to 72 time slots. Similarly, the choice of H depends on the accuracy of the grid model and grid load forecast. Finally, typical values for H^c can range from one to a few building time slots τ_B. In particular, we pick the contract start time t_{cs} and end time t_{ce} such that $t_{ce} - t_{cs} = H^c$, $t_{cs} \geq 1$ and $t_{cs} + 1 \leq t_{ce}$, in units of τ_B. In our simulations, we use $H^m = 24$ h, $H^c = \tau_B = 1$ h, $\tau_G = 1$ s and $H = 60$ s.

A summary of all the variables and parameters used to formalize the requirements and generate the MPC schemes is provided in Table 1 and Table 2. The overall set of requirements for the integrated energy management system can be formalized in terms of contracts as follows.

4.1 Building Contract

The building contract $\mathcal{C}_B$ can be expressed as follows.

Assumptions. Each building receives from the grid the vector of *electric energy prices* $\boldsymbol{\pi} = [\pi^e_t, \ldots, \pi^e_{t+H^m-1}]$ as well as the *prices of non-electric cooling and heating energy*, $\pi^{ne,c}$ and $\pi^{ne,h}$, respectively, every H^c time slots. Similarly, it receives

Table 1 Building and Building/Grid interface variables

Variable	Definition
	Building Requirements
H^m	Prediction horizon for the building MPC
$\mathscr{X}_t$	Set of permissible states at time t
$\mathscr{U}_t$	Set of permissible inputs at time t
	Building Model
τ_B	Sampling time for discretizing the building continuous dynamics
d_t	Disturbance at time t (e.g. outside temperature, occupancy, solar radiation)
T_t^{out}	Outside air temperature at time t
T^s	Supply air temperature exiting air handling unit (AHU)
COP_h	Coefficient of performance of heating system
COP_c	Coefficient of performance of cooling system
x_t	State of the system at time t
u_t	Input to the system at time t
w_t	Grid control uncertainty variable for the building control problem
P_f, P_h, P_c	Power consumption of fan, heating and cooling systems
	Building/Grid Interface Variables
H^c	Horizon (length) of the commercial contract
π_t^e	Per-unit price of electric energy at time t [\$/kWh]
$\pi^{ne,c}$	Per-unit price of non-electric cooling energy [\$/kWh]
$\pi^{ne,h}$	Per-unit price of non-electric heating energy [\$/kWh]
t_{cs}	Commercial contract start time
t_{ce}	Commercial contract end time
$\{\overline{\beta}_t, \underline{\beta}_t\}$	Reward paid from the grid to the building for upward flexibility ($\overline{\beta}_t$) and downward flexibility ($\underline{\beta}_t$) at time t [\$/kWh]
e_t^u	Upper envelope for safe air mass flow
e_t^l	Lower envelope for safe air mass flow
$\{\overline{\varphi}_t, \underline{\varphi}_t\}$	Upward ($\overline{\varphi}$) and downward ($\underline{\varphi}$) flexibility of the building at time t (in air flow)
$\{\overline{\psi}_t, \underline{\psi}_t\}$	Upward ($\overline{\psi}$), and downward ($\underline{\psi}$) flexibility of building at time t (in power)
$C_{\text{hvac}}(u_t, \pi_t^e)$	Total HVAC energy consumption cost at time t
$R(\Phi, \mathscr{B})$	Total reward from the grid to the building for flexibility

a pair of *rewards* vectors for providing upward flexibility ($\overline{\boldsymbol{\beta}}_t$) and downward flexibility ($\underline{\boldsymbol{\beta}}_t$). Finally, the building receives air flow control signals w_t from the grid, with the only assumption that they are bounded by the air flow flexibility as defined below, i.e. at each time t, we have $\underline{\varphi}_t \leq w_t \leq \overline{\varphi}_t$.

Guarantees. The building must satisfy a set of *temperature requirements*, expressed as predicates on the building states of the form "x_{t+k} should be in $\mathscr{X}_{t+k}$ for all times $t+k$ where $k \in \{1,\ldots,H^m\}$". Similarly, *air mass flow requirements* can be formalized as a set of constraints $\mathscr{U}_{t+k}$ on the building inputs for all $k \in \{0,\ldots,H^m-1\}$, where

$$\mathscr{X}_t := \{x \mid \underline{T}_t \leq x \leq \overline{T}_t\} \tag{5}$$

$$\mathscr{U}_t := \{u \mid \underline{U}_t \leq u \leq \overline{U}_t\}, \tag{6}$$

Table 2 Power system variables

Variable	Definition
H	Prediction horizon for the grid MPC
λ	Bound on the rate of change of the power supplied by the buildings
τ_G	Sampling time for discretizing the grid continuous dynamics
P_M	Mechanical power input
P_M^o	Desired real power generation
P_G	Generated real electric power
δP_G	Increase in demand (at rated generator MVA)
V_t	Terminal voltage
P_D	Load (power demand)
δP_D	Input disturbance due to load changes
δP_C	Speed changer position feedback control signal
ω	Angular speed and frequency
ω_o	Rated (desired) frequency
D	Damping coefficient. Range: 0.01 - 0.1
M	Machine inertia constant. Range: 100 - 1000 [MW s]
R	Speed regulation constant. Range: 0.05 [p.u.]
T_i	Time constant for power system components. Range: {0,0.01-10} [s]
K_i	Fraction of total mechanical power outputs associated with different operating points of the turbine. Range: {0,0.1-1}

$\underline{T}_t$ and $\overline{T}_t$ are the upper and lower temperature limits, and $\overline{U}_t$ and $\underline{U}_t$ are the upper and lower feasible air mass flow rates at time t.

We say that the building offers a *flexibility* $\Psi := \{\underline{\boldsymbol{\psi}}, \overline{\boldsymbol{\psi}}\}$ in fan power or equivalently a flexibility $\Phi := \{\underline{\boldsymbol{\varphi}}, \overline{\boldsymbol{\varphi}}\}$ in air mass flow, including downward flexibility $\underline{\boldsymbol{\varphi}}$ and upward flexibility $\overline{\boldsymbol{\varphi}}$, from the contract start time t_{cs} to the contract end time t_{ce} if there exist two trajectories $\mathbf{e^l} = \mathbf{u} + \underline{\boldsymbol{\varphi}}$ and $\mathbf{e^u} = \mathbf{u} + \overline{\boldsymbol{\varphi}}$, that satisfy:

$$\underline{\varphi}_k \leq 0, \quad \overline{\varphi_k} \geq 0 \qquad \forall k \in \{t_{cs}, \ldots, t_{ce}\} \tag{7}$$

$$f(x_k, u_k + \underline{\varphi}_k, d_k) \in \mathscr{X}_{k+1} \qquad \forall k \in \{t_{cs} - 1, \ldots, t_{ce} - 1\} \tag{8}$$

$$f(x_k, u_k + \overline{\varphi}_k, d_k) \in \mathscr{X}_{k+1} \qquad \forall k \in \{t_{cs} - 1, \ldots, t_{ce} - 1\} \tag{9}$$

$$u_k + \underline{\varphi}_k \in \mathscr{U}_k \qquad \forall k \in \{t_{cs}, \ldots, t_{ce}\} \tag{10}$$

$$u_k + \overline{\varphi}_k \in \mathscr{U}_k \qquad \forall k \in \{t_{cs}, \ldots, t_{ce}\} \tag{11}$$

where f captures the building dynamics, and d_k is an estimate of unmodelled disturbances, as detailed in Section 5.1. If the BEMS declares $\underline{\boldsymbol{\varphi}}$ and $\overline{\boldsymbol{\varphi}}$, then the utility can choose any fan power (and consequently air flow $\hat{u}_k$) for all time steps $t_{cs} \leq k \leq t_{ce}$ as long as $u_k^* + \underline{\varphi}_k \leq \hat{u}_k \leq u_k^* + \overline{\varphi}_k$, where u_k^* is the *baseline* air mass flow. Hence, we "center" the flexibility around u^*.

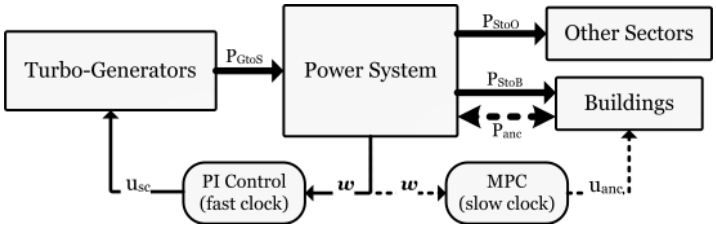

Fig. 5 Schematic of the power system showing its interconnections with the turbo-generators, the building subsystem, and other sectors, along with the control architecture. Thick arrows represent the power flow while thin arrows represent frequency and control signals. Dashed arrows indicate the additional signals and power flows proposed in this work, on top of the state-of-the-art AGC.

4.2 Grid Contract

To express the grid contract $\mathcal{C}_G$, we refer to the grid system architecture in Fig. 5, where commercial buildings are abstracted as a load for the power system.

Assumptions. Buildings can provide both positive and negative power flow to the grid for frequency regulation purposes. When there is a power deficit, buildings will temporarily use less power, and when there is as surplus of power, they will temporarily use the extra power.

To formalize such a behavior at the grid level, we abstract the building subsystem in terms of: (i) aggregate baseline power demand P_D, and (ii) aggregate flexibility for ancillary services δP_{anc}. In practice, the grid operator predicts both the long-term power demand and its short-term deviations from historical data (e.g. weather patterns) by using machine learning algorithms. Therefore, in this formulation, we assume that the power demand P_D, and any variation δP_D, are known parameters, internal to the power system model. Similarly, the upper bound on the ramping rate of the ancillary service power $\lambda > 0$ is treated as a constant parameter of the power system model. Then, the only assumptions of the grid operator on the characteristics of the ancillary service signal from the buildings can be expressed in terms of *maximum capacity* $\max(\delta P_{\text{anc}_k}) > 0$, and a *minimum capacity* $\min(\delta P_{\text{anc}_k}) < 0$ at each time k. Such bounds can be directly derived by the flexibility declared by the buildings at the beginning of the commercial contract. Finally, we assume that u_{sc} in Fig. 5 is also constant, since it is regulated by a local PI controller external to the G-MPC problem.

Guarantees. The utility guarantees that its power consumption signals to the buildings will be within the assumed maximum and minimum capacity boundaries above. Moreover, the recent Federal Energy Regulatory Commission (FERC) Order 755 requires scheduling coordinators to procure and compensate more for regulation resources with faster ramping rates. This ramping rate constraint can be formalized by requiring that $|\Delta P_{\text{anc}}| = |\delta P_{\text{anc}_{k+1}} - \delta P_{\text{anc}_k}| \leq \lambda$ holds at each time step k, i.e. the *rate of change* of the power supplied by the buildings must be guaranteed to be limited by λ.

5 System Model Library

We describe the models of the different components of the grid, starting with the building subsystem. These models will be used together with the contracts in Section 4 to formulate the MPC optimization problems.

5.1 Building Model

We equip the building component with both a *behavioral model*, capturing the system dynamics, and a non-functional model, capturing the electric power consumption as a function of the air mass flow u_t and the outer temperature T_t^{out} at time t. By assuming a discretization step τ_B as in Section 4, the building dynamics are regulated by a difference equation of the form

$$x_{t+1} = f(x_t, u_t, d_t) \tag{12}$$

where x_t represent the system state, i.e. the temperatures of different rooms or zones in the building, u_t is the air mass flow to the thermal zones, and d_t is an estimate of the unmodelled disturbances, e.g. outside temperature or building occupancy [17]. The function $f(.)$ is generally non-linear, which makes it more difficult to handle in an optimization framework. Therefore, we adopt a linearized expression of the state update equation using the forward Euler integration formula with time-step τ_B as:

$$x_{t+1} = Ax_t + Bu_t + Ed_t. \tag{13}$$

The building HVAC *power consumption* is the summation of fan power, cooling power and heating power. With the assumption of no recirculation of air and constant air mass flow, the three contributions can be calculated as follows:

$$P_f(u_t) = c_1 u_t^3 + c_2 u_t^2 + c_3 u_t + c_4 \tag{14}$$

$$P_h(u_t, T_t^{out}) = c_p u_t (T^s - T_t^{out})/COP_{\mathrm{h}} \tag{15}$$

$$P_c(u_t, T_t^{out}) = c_p u_t (T_t^{out} - T^s)/COP_{\mathrm{c}}, \tag{16}$$

where T^{out} is the outside air temperature, constants c_{1-4} are fan parameters, c_p is the specific heat of air, COP_{h} and COP_{c} are the performance coefficients for the heating system and the cooling system, respectively, and the supply air temperature T^s is considered constant. To move the coolant fluid around, heating and cooling systems use pumps which consume electric power. However, we assume that electric power consumption of pumps is negligible compared to the non-electric heating and cooling powers of these systems [18, 16].

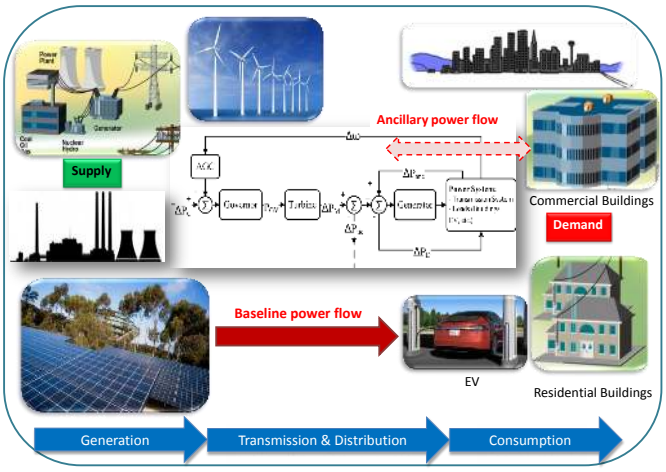

Fig. 6 Block diagram of power system and its relation to governor, turbine, generator, and the AGC signal for each control area.

5.2 Grid Model

The power system model, based on [8, 34, 13], consists of a governor, turbine, and generator, interconnected as shown in the block diagram of Fig. 6. Electric power is generated by the turbo-generators, and is fed to the power system. The power system transmits and distributes the power to the end users. δP_C is a control input which acts against increase or decrease in the power demand to regulate the system frequency. δP_D denotes the fluctuations in the power demand, which are here considered as an exogenous input (disturbance), while the aggregate flexibility from all the buildings participating in the SF program is lumped into δP_{anc}. As discussed in Section 4.2, we assume that both the power demand P_D and its variations δP_D are known from historical data. The variables of the power system are listed in Table 2 together with a short description for each of them. Our model relies on the following simplifying assumptions:

- The resistance of the transmission lines is ignored;
- The transmission line between areas i and j is characterized by a reactance $X_{\text{tie}_{ij}}$;
- Reactive power flows are ignored;
- The voltage V_i of bus i is considered constant.

At steady state, we have: $\omega = \omega_o$, $V_t = V_t^o$, and $P_M = P_G = P_M^o$, where ω_o, V_t^o, and P_M^o are the nominal values for rated frequency, terminal voltage and mechanical power input. We are interested in modeling the incremental changes with respect to the steady-state condition.

Governor. The overall input-output transfer function of the governor is given by

$$F_{Gov}(s) = \frac{(1+sT_2)}{(1+sT_1)(1+sT_3)}. \tag{17}$$

Typical values for the time constants depend on whether governors are mechanical-hydraulic or electro-hydraulic, with or without steam feedback [8].

Turbine. The input-output transfer function of a turbine is given by

$$\frac{\delta P_M}{P_{GV}} = K_1F_1 + K_3F_1F_2 + K_5F_1F_2F_3 + K_7F_1F_2F_3F_4, \tag{18}$$

where F_1, F_2, F_3, and F_4 are transfer functions corresponding to steam chest, piping system, re-heaters, and cross-over mechanisms, respectively, and are given by

$$F_1(s) = \frac{1}{1+sT_4}, \qquad F_2(s) = \frac{1}{1+sT_5}, \tag{19}$$

$$F_3(s) = \frac{1}{1+sT_6}, \qquad F_4(s) = \frac{1}{1+sT_7}. \tag{20}$$

The basic time constant associated with steam turbines is T_4 (*steam chest*). For non-reheat steam turbines, this is the only needed time constant. The coefficients K_1, K_3, K_5, and K_7 represent fractions of the total mechanical power outputs associated with *very high*, *high*, *intermediate*, and *low* pressure components, respectively. Typical values of the steam turbine time constants and fractions are reported in [8].

Generator. The dynamics of the generator is given by the following transfer function

$$F_{Gen} = \frac{1}{D+sM}, \tag{21}$$

where constants D and M represent the damping coefficient and the inertia of the governor, respectively.

Two Area System Model. The components above can be interconnected to generate a model for the system in Fig. 4, including two areas connected by a tie line with reactance X_{tie}. P_{tie} is the power flow on the tie line from area 1 to area 2. A positive δP_{tie} represents an increase in power transfer from area 1 to area 2. This in effect is equivalent to increasing the load of area 1 and decreasing the load of area 2. The model can be directly extended to n areas, under the assumption that non-negligible power transfers can only occur between area i and its nearest neighbors $i-1$ and $i+1$. In the two-area model, the superscripts refer to the control area ($i, j \in \{1,2\}$) and the subscripts index the state in each area:

$$\frac{dx_1^i}{dt} = \frac{(-D^i x_1^i + \delta P_M^i - \delta P_D^i - \delta P_{\text{tie}}^{ij} + \delta P_{\text{anc}}^i)}{M_x^i} \tag{22a}$$

$$\frac{dx_2^i}{dt} = \frac{(x_3^i - x_2^i)}{T_7^i} \tag{22b}$$

$$\frac{dx_3^i}{dt} = \frac{(x_4^i - x_3^i)}{T_6^i} \tag{22c}$$

$$\frac{dx_4^i}{dt} = \frac{(x_5^i - x_4^i)}{T_5^i} \tag{22d}$$

$$\frac{dx_5^i}{dt} = \frac{(P_{GV}^i - x_5^i)}{T_4^i} \tag{22e}$$

$$\frac{dx_6^i}{dt} = \frac{(x_7^i - x_6^i)}{T_3^i} \tag{22f}$$

$$\frac{dx_7^i}{dt} = \frac{(-x_7^i + \delta P_C^i - x_1^i/R^i)}{T_1^i} \tag{22g}$$

where

$$\delta P_M^i = K_1^i x_5^i + K_3^i x_4^i + K_5^i x_3^i + K_7^i x_2^i \tag{23}$$

$$P_{GV}^i = (1 - T_2/T_3)x_6^i + (T_2/T_3)x_7^i \tag{24}$$

In (22), the first state variable represents the frequency increment, $x_1^i = \delta\omega_i$. The differential equations for the seven state variables are derived using the mathematical models in (17)–(21). When the time constant representing the system pole is zero, the corresponding differential equation becomes an algebraic equation. For instance, when $T_5 = 0$, the equation $\frac{dx_4^i}{dt} = 1/T_5(x_5^i - x_4^i)$ turns into $x_5^i = x_4^i = 0$.

The real power transferred from bus i to bus j can be approximated as $P_{\text{tie}}^{ij} \approx V_i V_j b_{ij} cos(\theta_i - \theta_j)$. Since here we are concerned with incremental changes in all variables, the incremental change in P_{tie}^{ij} is given by $\delta P_{\text{tie}}^{ij} = \nu_{ij}(\theta_i - \theta_j)$ where at the nominal operating points, θ_i^o, $i = 1,2$, the *transmission line stiffness* coefficient ν_{ij} is given by

$$\nu_{ij} = -V_i V_j b_{ij} cos(\theta_i^o - \theta_j^o). \tag{25}$$

In terms of the incremental state variables used, we have:

$$\delta P_{\text{tie}}^i = \sum_{j=1}^{n} \nu_{ij}(x_8^i - x_8^j), \tag{26}$$

where the state variable x_8^i is the integral of the frequency increment of area i, i.e.,

$$\frac{dx_8^i}{dt} = x_1^i. \tag{27}$$

The state space model (22)-(27) can be written in compact form as follows:

$$\frac{dx(t)}{dt} = A'x(t) + B'_1 u^{\text{sc}}(t) + B'_2 u^{\text{anc}}(t) + E'd(t). \tag{28}$$

States are stored in x, input signals to the speed changers are $u^{\text{sc}} = [\delta P_{C_1} \ \delta P_{C_2}]^T$, the ancillary inputs from the buildings are $u^{\text{anc}} = [\delta P_{\text{anc}_1} \ \delta P_{\text{anc}_2}]^T$, and the exogenous inputs (disturbance), modeling the variations in demand are denoted by $d = [\delta P_{D_1} \ \delta P_{D_2}]^T$.

We discretize the state space dynamics using the forward Euler scheme. We show the result on the equation for dx_1^i/dt. The discretized dynamics for the other states can be obtained in a similar way. At time t_n we approximate the derivative of x_1^i by

$$\frac{dx_1^i(t_n)}{dt} \approx \frac{x_1^i(t_n + \delta t) - x_1^i(t_n)}{\delta t} \tag{29}$$

Hence the discretized version of (22a) is

$$x_1^i(t_{n+1}) \quad = \quad \left(1 - \frac{D_x^i \delta t}{M_x^i}\right) x_1^i(t_n) \quad + \quad \frac{\delta t}{M_x^i}\left[\delta P_M^i - \delta P_D^i - \delta P_{\text{tie}}^{ij} + \delta P_{\text{anc}}^i\right],$$

where $t_{n+1} = t_n + \delta t$ and $\delta t = \tau_G$ is the discretization time step. The discrete-time state-space model is obtained as

$$x_{n+1} = Ax_n + B_1 u_n^{\text{sc}} + B_2 u_n^{\text{anc}} + Ed_n. \tag{30}$$

We use this state update equation for the G-MPC problem in Section 6.3.

Automatic Generation Control. The AGC is the main control function of a utility's energy control section. The purpose of an AGC is to track the load variations while maintaining the system frequency, net tie-line interchanges, and optimal generation level close to scheduled values [8]. This function is referred to as Load-Frequency Control. A secondary objective is to distribute the required change in generation among units to minimize operating costs [13]. In the case where several areas are interconnected, each will perform its own AGC independent of the others.

In the classical AGC, a simple PI control is utilized to regulate the frequency of the grid. The Area Control Error (ACE) is defined as

$$ACE^i = \delta P_{tie}^i + \beta^i x_1^i, \tag{31}$$

where $\delta P_{\text{tie}}^i = P_{\text{tie}}^i - P_{\text{tie,scheduled}}^i$, and β^i is the bias coefficient of area i. The standard industry practice is to set the bias β^i at the so-called Area Frequency Response Characteristic which is defined as $\beta^i = D^i + 1/R^i$. The integral of the ACE is then used to construct the speed changer position feedback control signal δP_C^i. A new state x_9^i is then defined as

$$\frac{dx_9^i}{dt} = ACE^i. \tag{32}$$

Consequently the control input δP_C^i is given by

$$\delta P_C^i = -K^i x_9^i, \tag{33}$$

where K^i is the feedback gain. The MPC scheme proposed in Sec. 6.3 controls the available ancillary service from commercial buildings to improve on the classical AGC practice. This optimization-based control framework is utilized as a *higher-level* control in a "hierarchical" fashion on top of the *low-level* classical AGC control, as visualized in Fig. 5.

6 Energy Management System Optimal Control

We first describe the communication protocol followed by the grid and the building subsystem over time. Then, based on this protocol, we show how the building and the grid MPC schemes can be generated, respectively, from the contracts in Section 4 and the models in Section 5.

6.1 Grid-Building Communication Protocol

The grid operator and the building subsystem communicate as follows:

1. Buildings and grid agree upon the length H^c of the commercial contract.
2. The utility declares $\boldsymbol{\pi} = [\pi_0^e, \ldots, \pi_{H^m-1}^e]$, the vector of prices of electric energy per unit step, the vector $\underline{\boldsymbol{\beta}} = [\underline{\beta}_0, \ldots, \underline{\beta}_{H^m-1}]$ of rewards for downward flexibility, and the vector $\overline{\boldsymbol{\beta}} = [\overline{\beta}_0, \ldots, \overline{\beta}_{H^m-1}]$ of reward for upward flexibility. If the utility is not willing to commit to the flexibility rates for the time span beyond the next, immediate contract period, e.g. $[\underline{\beta}_{H^c+1}, \ldots, \underline{\beta}_{H^m-1}]$, and $[\overline{\beta}_{H^c+1}, \ldots, \overline{\beta}_{H^m-1}]$, each building operator can obtain an estimate of these values from historical data. The same can be stated for the prices of electric energy beyond the next contract period $[\pi_{H^c+1}^e, \ldots, \pi_{H^m-1}^e]$.
3. Each building operator computes the baseline air mass flow u_k^* and the two envelopes e_k^l and e_k^u, for the time frame $k \in \{0, 1, \ldots, H^m - 1\}$, by solving the B-MPC, and declares the envelope $P_f(\mathbf{e^l})$, $P_f(\mathbf{e^u})$ and the baseline $P_f(\mathbf{u}^*)$ power consumption profiles. The B-MPC can be safely solved *independently* of the power control signals received from the grid, since the buildings do not make any assumption on such signals, but the fact that they are confined within the declared power envelopes. It is then possible for the buildings to separately minimize their cost *for all admissible values* of the grid control signals. Such a robust

control problem is key to compositional design since it breaks the circularity between the B-MPC and the G-MPC problems.

4. The grid operator aggregates the envelope profiles received from the buildings and repeatedly solves the G-MPC, to obtain the control signals u_j^{anc} for $j \in \{0,1,\ldots,H\}$. These control signals are the aggregation of all the ancillary service powers provided by the buildings. At each time step, the grid operator will disaggregate δP_{anc_i} for each control area i into N pieces proportional to the declared flexibility of each building, such that $\delta P_{\text{anc}_i} = \sum_{k=1}^{N} s_i^k$, where N is the number of buildings participating in the ancillary service program.
5. During the next H^c time slots, the grid operator will send signals s_j^k in each control area, such that $P_f(e_j^{l,k}) \leq s_j^k \leq P_f(e_j^{u,k})$ and the building operator k has to track the signals, i.e., has to consume power equal to s_j^k in time slot j. The flexibility signal s_j^k may arrive as frequently as every few seconds, as mentioned earlier.

6.2 Design of the Building Optimal Control Scheme

Let H^m be the prediction horizon of the B-MPC in terms of building time slots, selected as discussed in Section 4. We design the B-MPC scheme to minimize the building economic cost in terms of baseline power consumption and flexibility over H^m, while satisfying the building constraints and the occupants' comfort requirements as encoded by the building contract $\mathcal{C}_B$ in Section 4.1. At each time t, the predictive controller solves an optimization problem to compute the baseline air mass flow $\mathbf{u}_t = [u_t, \ldots, u_{t+H^m-1}]$ to the thermal zones of the building, and the downward and upward mass flow flexibility vectors $\underline{\boldsymbol{\varphi}}_{t+1} = [\underline{\varphi}_{t+1}, \ldots, \underline{\varphi}_{t+H^m}]$ and $\overline{\boldsymbol{\varphi}}_{t+1} = [\overline{\varphi}_{t+1}, \ldots, \overline{\varphi}_{t+H^m}]$.

The inputs to the optimization problem are the *initial state* x_t (zone temperatures), the set of *electric* energy prices per time slot $\{\pi_t^e, \ldots, \pi_{t+H^m-1}^e\}$, the *non-electric* energy prices, such as gas price for heating $\pi^{ne,h}$ and cooling $\pi^{ne,c}$ (which are considered time-invariant), the prediction of the outside temperature and inside heat generation d_t. The cost function consists of the *cost* for the baseline HVAC power consumption, C_{hvac}, minus the *reward* for the flexibility R, computed as follows:

$$R(\Phi, \mathscr{B}) = \overline{\boldsymbol{\beta}}^T \overline{\boldsymbol{\psi}}(\mathbf{u}, \overline{\boldsymbol{\varphi}}) + \underline{\boldsymbol{\beta}}^T \underline{\boldsymbol{\psi}}(\mathbf{u}, \underline{\boldsymbol{\varphi}}) \tag{34}$$

where $\mathscr{B} := \{\underline{\boldsymbol{\beta}}, \overline{\boldsymbol{\beta}}\}$, with $\underline{\boldsymbol{\beta}}$ and $\overline{\boldsymbol{\beta}}$ as in Section 6.1, $\Phi := \{\underline{\boldsymbol{\varphi}}, \overline{\boldsymbol{\varphi}}\}$, and $\underline{\boldsymbol{\psi}}(.)$ and $\overline{\boldsymbol{\psi}}(.)$ are given, at each time step k, by

$$\overline{\psi}(u_k, e_k^u) \triangleq P_f(e_k^u) - P_f(u_k) \tag{35a}$$

$$\underline{\psi}(u_k, e_k^l) \triangleq P_f(u_k) - P_f(e_k^l) \tag{35b}$$

in which $\mathbf{e}^{\mathbf{l}} = \mathbf{u} + \underline{\boldsymbol{\varphi}}$ and $\mathbf{e}^{\mathbf{u}} = \mathbf{u} + \overline{\boldsymbol{\varphi}}$. The total HVAC power consumption cost $C_{\text{hvac}}(u_t, \pi_t)$ is the summation of fan power, cooling power and heating power costs, given by:

$$C_{\text{hvac}}(u_t, \pi_t) = \pi_t^e P_f(u_t) + \pi^{ne,c} P_c(u_t, T_t^{out}) + \pi^{ne,h} P_h(u_t, T_t^{out}) \tag{36}$$

where T^{out} is the outside air temperature, and the three power contributions are calculated based on the building model in Section 5.1.

As discussed in Section 6.1, the building operator aims to minimize the economic cost in the worst-case scenario for the grid signals (environment) following a game-theoretic approach. The result is a robust *min-max* optimization problem. The inner maximization problem derives the worst-case scenario cost and constraints, while the outer minimization problem solves for its arguments $(\mathbf{u}_t, \Phi_{t+1})$, while guaranteeing that the constraints are satisfied for all the values of the uncertain mass flow signals $\mathbf{w}$ from the grid, as long as they are within the range allowed by the building flexibility. Therefore, at time t the building operator solves:

$$\min_{\mathbf{u}_t, \Phi_{t+1}} \max_{\mathbf{w}_t} \sum_{k=0}^{H^m-1} C_{\text{hvac}}(u_{t+k}, \pi_{t+k}) - R(\Phi_{t+k+1}, \mathscr{B}_{t+k+1}) \tag{37a}$$

$$\text{subject to:} \quad x_{t+k+1} = f(x_{t+k}, u_{t+k} + w_{t+k}, d_{t+k}), \quad \forall\, k \in \{0, \ldots, H^m - 1\} \tag{37b}$$

$$\forall\, w_t \text{ s.t. } \underline{\varphi}_t \leq w_t \leq \overline{\varphi}_t \tag{37c}$$

$$\forall\, w_{t+k} \text{ s.t. } \underline{\varphi}_{t+k} \leq w_{t+k} \leq \overline{\varphi}_{t+k}, \quad \forall\, k \in \{1, \ldots, H^m - 1\} \tag{37d}$$

$$\overline{\varphi}_{t+k} \geq 0, \ \forall\, k \in \{1, \ldots, H^m - 1\} \tag{37e}$$

$$\underline{\varphi}_{t+k} \leq 0, \ \forall\, k \in \{1, \ldots, H^m - 1\} \tag{37f}$$

$$x_{t+k} \in \mathscr{X}_{t+k}, \ \forall\, k \in \{1, \ldots, H^m\} \tag{37g}$$

$$u_{t+k} + w_{t+k} \in \mathscr{U}_{t+k}, \ \forall\, k \in \{0, \ldots, H^m - 1\} \tag{37h}$$

We observe that $\underline{\varphi}_t$ and $\overline{\varphi}_t$ are computed in the previous time step and are constant values in (37), while $\underline{\varphi}_{t+k}$ and $\overline{\varphi}_{t+k}$ for $k \in \{1, \ldots, H^m - 1\}$ are optimization variables and will be computed in the current time step by solving the optimal control problem. Therefore, B-MPC computes the future flexibility profile (starting from the next time step), based on the current flexibility. For the very first time step, we assume $\underline{\varphi}_0 = \overline{\varphi}_0 = 0$.

The inner optimization problem can be solved analytically. In fact, according to a fundamental theorem on convex functions [6], if a convex function attains a maximum over a closed convex set, then the maximum is achieved at some extreme point of the set. When the building state update equation are linearized as described in Section 5.1, the feasible set for states (the temperature of the rooms in the building) and controls (air mass flow into the thermal zones) is closed and convex, being an intersection of closed half-spaces. The objective function is also convex in $\mathbf{w}_t$, since $\mathbf{w}_t$ does not appear in the cost function. Hence, the min-max problem (37) is equivalent to:

$$\min_{\mathbf{u}_t, \Phi_{t+1}} \quad \sum_{k=0}^{H^m-1} C_{\text{hvac}}(u_{t+k}, \pi_{t+k}) - R(\Phi_{t+k+1}, \mathscr{B}_{t+k+1}) \tag{38a}$$

$$\text{s. t.:} \quad \underline{x}_{t+k+1} = f(x_{t+k}, u_{t+k} + \underline{\varphi}_{t+k}, d_{t+k}) \quad \forall k \in \{0,...,H^m-1\} \tag{38b}$$

$$\overline{x}_{t+k+1} = f(x_{t+k}, u_{t+k} + \overline{\varphi}_{t+k}, d_{t+k}) \quad \forall k \in \{0,...,H^m-1\} \tag{38c}$$

$$\overline{\varphi}_{t+k} \geq 0, \quad \forall k \in \{1,...,H^m-1\} \tag{38d}$$

$$\underline{\varphi}_{t+k} \leq 0, \quad \forall k \in \{1,...,H^m-1\} \tag{38e}$$

$$\underline{x}_{t+k} \in \mathscr{X}_{t+k} \quad \forall k \in \{1,...,H^m-1\} \tag{38f}$$

$$\overline{x}_{t+k} \in \mathscr{X}_{t+k} \quad \forall k \in \{1,...,H^m-1\} \tag{38g}$$

$$u_{t+k} + \underline{\varphi}_{t+k} \in \mathscr{U}_{t+k} \ \forall k \in \{0,...,H^m-1\} \tag{38h}$$

$$u_{t+k} + \overline{\varphi}_{t+k} \in \mathscr{U}_{t+k} \ \forall k \in \{0,...,H^m-1\} \tag{38i}$$

The result of (38) is the nominal power consumption u^*_{t+k} and the maximum available flexibility Φ^*_{t+k+1}, for $k \in \{0,...,H^m-1\}$. The building declares u^*_{t+k} and Φ^*_{t+k+1} for $k \in \{0,...,H^c-1\}$ to the utility. After H^c time slots, the BEMS collects the updated parameters such as new measurements and disturbance predictions, sets up the new MPC algorithm for $k \in \{H^c, H^c+1, \ldots, H^c+H^m-1\}$, solves the new MPC for this time frame and uses only the first H^c values of baseline power consumption and flexibility, i.e. for $k = H^c, \ldots, 2H^c-1$, and this process repeats.

6.3 Design of the Grid Optimal Control Scheme

We design the G-MPC problem by combining assumptions and guarantees under the responsibility of the grid operator, as expressed by the contract $\mathcal{C}_G$ in Section 4.2. Let $U^{\text{anc}}_k = (u^{\text{anc}}_k, u^{\text{anc}}_{k+1}, \ \ldots, u^{\text{anc}}_{k+H-1})$ the trajectory of the power control signal of the grid to the buildings for time steps from k to $k+H-1$, where H is the prediction horizon of G-MPC. We aim to minimize the ℓ_2 norm of the *ACE* signal in areas $i = 1,2,...,n$, by exploiting the ancillary service available from buildings, taking into account the system dynamics and constraints. More formally, at each time step k, we solve:

$$\min_{U^{\text{anc}}_k} \quad \sum_{i=1}^{n} \sum_{j=0}^{H-1} (\text{ACE}^i_{k+j})^2 \tag{39}$$

$$\text{s.t.} \quad x_{k+j+1} = Ax_{k+j} + B_2 u^{\text{anc}}_{k+j} + Ed_{k+j}$$

$$\underline{\mu}_{k+j} \leq u^{\text{anc}}_{k+j} \leq \overline{\mu}_{k+j}$$

$$|u^{\text{anc}}_{k+j+1} - u^{\text{anc}}_{k+j}| \leq \lambda_{k+j}$$

where, for each area i, $\overline{\mu}_{k+j}(i) = \sum_{m=1}^{N} P_f(e^{u,m,i}_{k+j})$ and $\underline{\mu}_{k+j}(i) = \sum_{m=1}^{N} P_f(e^{l,m,i}_{k+j})$. All the constraints of problem (39) should hold for $j = 0,1,\ldots,H-1$. The constraints

of the optimization problem are $\overline{\mu}_{k+j} > 0$ for the maximum positive power and $\underline{\mu}_{k+j} < 0$ for maximum negative power provided by the set of buildings in each area. Here, "positive" and "negative" refer to the flow of power from generation to consumption. These values are computed by the buildings and sent to the grid operator periodically, as detailed in Section 6.1. λ_{k+j} is the maximum limit on the rate of change of ancillary service provided by the buildings. Based on the assumptions in Section 4.1, deviations in the power demand from the buildings are known and lumped into the signal d_{k+j}. A robust version of the G-MPC problem to deal with uncertainties in the power demand, following a similar approach as in the B-MPC problem, will be object of future work. Finally, we do not incorporate $B_1 u_k^{sc}$ in the state-space model, since we assume that u^{sc} is constant and regulated by the local PI controller.

7 Simulation Results

To validate the proposed methodology, we simulated the control algorithms in Section 6 by using the building model in [16, 17], developed and validated against historical data. For rapid prototyping, we used YALMIP [14] as an interface to back-end optimization solvers. The non-linear optimization problem in B-MPC was solved using IPOPT [35], while CPLEX [3] was used to solve the quadratic program generated by G-MPC.

7.1 Validation of the B-MPC Algorithm

Different reward rates have been considered for upward and downward flexibility at each time step, as shown at the bottom of Fig. 7, under the constraint that downward flexibility is rewarded more than upward flexibility, i.e. $\underline{\beta} > \overline{\beta}$ for most of the day. We performed simulations with a sampling time of 1 hour and a prediction horizon of $H^m = 24$ h. On a 4-core 2.67-GHz Intel processor with 3.86 GB of memory, the mean and standard deviation of solver times were 8.9 s and 5.3 s, respectively. Fig. 7 shows the results when ancillary signals are received from the grid every minute: no building constraint (e.g. temperature comfort zone) is violated for arbitrary values of the fan speed enforced by the grid, as long as the fan power consumption is within the safe envelope calculated by B-MPC. The maximum flexibility (100%) is obtained when the room temperature is far from the boundaries of the comfort zone. The flexibility decreases as the temperature of the room approaches the comfort zone boundary, and reaches its minimum (about 0-15%) when the room temperature is close to the boundaries of the comfort zone, and the reward is small. Fig. 7 shows that the control strategy in (38) can indeed offer HVAC energy consumption flexibility via proper incentives. Table 3 lists the parameters used in our simulations.

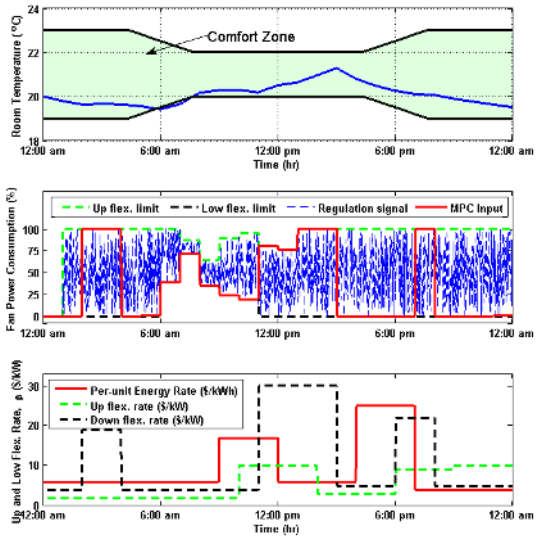

Fig. 7 Per-unit energy rate, upward and downward flexibility reward (bottom), building flexibility (middle), room temperature profile (top). Flexibility signals are sent every minute from the grid.

Table 3 Simulation parameters used for B-MPC validation.

Parameter	Value
c_1	-6.06×10^{-13} $[CFM^{-3}]$
c_2	6.73×10^{-8} $[CFM^{-2}]$
c_3	-1.2×10^{-3} $[CFM^{-1}]$
c_4	59.2
COP_h	3
COP_c	2
c_p	1.0 [kJ/(kg·K)]

7.2 Validation of the G-MPC Algorithm

We consider two interconnected control areas with model parameters as in Table 4, and with inter-area stiffness coefficient $\nu = 1.0$ p.u. The main generation unit for area 1 is a non-reheat turbo generator (TG) system while the main generation unit for area 2 is a hydro TG system. Some metrics such as root mean square (rms) values of frequency and ACE signal are considered to compare the performance of the proposed controller with respect to a traditional scheme. To simplify, we use time-invariant bounds for the maximum and minimum ancillary power $\overline{\mu}_k = \underline{\mu}_k = \mu$ and maximum rate of change of ancillary power λ in the following simulations.

We performed simulations for a time horizon of 100 s and a sampling time $\tau_G = 1$ s. The mean and standard deviation solver times were 0.02 s and 0.005 s,

Table 4 Simulation parameters used for G-MPC validation.

Control Area	Parameters
Area 1	$T_1 = 0.1$, $T_3 = 0.1$, $T_4 = 1.0$
	$K_1 = 1.0$
	$M = 132.6$ [MW s]
	$D = 0.0265$ [p.u.]
Area 2	$T_1 = 0.2$, $T_3 = 0.3$, $T_4 = 0.1$, $T_5 = 0.5$
	$K_1 = 0.2$, $K_3 = 3$
	$M = 663.13$ [MW s]
	$D = 0.1325$ [p.u.]

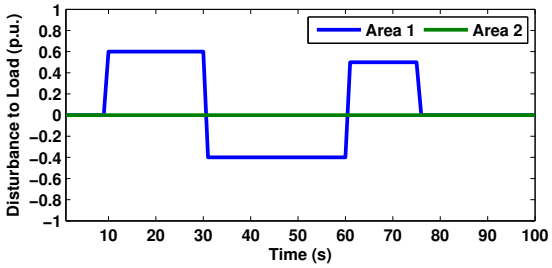

Fig. 8 Load disturbance signal.

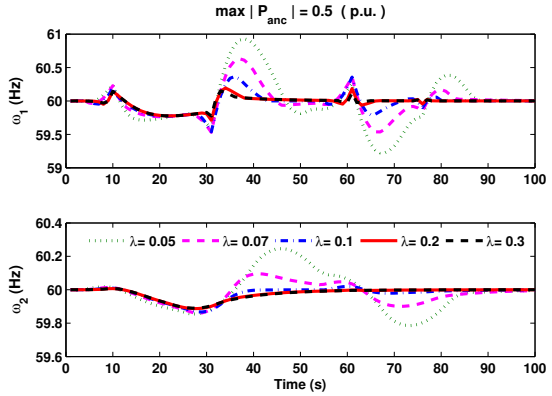

Fig. 9 Frequency of areas 1 and 2 in response to the load disturbance. The prediction horizon is $H = 10$ and the maximum ancillary power is $\max(P_{\text{anc}}) = 0.5$ p.u. Results relate to different values of rate-of-change of ancillary power ($\max|\Delta P_{\text{anc}}| = \lambda$).

respectively, for a prediction horizon of $H = 10$ s. We consider a disturbance signal in the load of area 1, and no disturbance in the load of area 2, as shown in Fig. 8. We assess the performance of the proposed controller considering the following scenarios.

Scenario 1. The maximum ancillary service available in each area is 0.5 per unit (p.u.) of power. We consider a prediction horizon of $H = 10$ time steps. As shown in

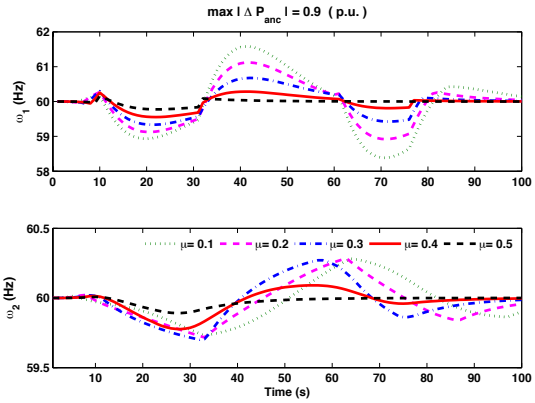

Fig. 10 Frequency trajectories in control areas 1 and 2 ($H = 10$, $\max(\Delta P_{\text{anc}}) = 0.9$ p.u., $\max|P_{\text{anc}}| = \mu$).

Fig. 9, by increasing the maximum rate of change of ancillary power (ramping rate for generation units) the resulting frequency deviation decreases. A ramping rate of $\lambda = 0.05$ (p.u./s) is associated with large power generator size, a ramping rate of $\lambda = 0.1$ (p.u./s) is associated with smaller size generators. High ramping rates, such as $\lambda = 0.3$ (p.u./s), are associated with fast ancillary services such as the one provided by building HVAC system fans.

Scenario 2. To investigate the impact of the maximum available ancillary power constraint on the frequency deviations in each area, we relax the constraint on the ramping rate of the ancillary services, by selecting a higher bound $\lambda = 0.9$ (p.u./s). As shown in Fig. 10, by increasing the maximum available ancillary power μ the frequency deviation decreases, thus showing the effectiveness of our control strategy. The disturbance in the load of area 1 affects both the interconnected areas; however, the frequency change in area 1 is larger than the one in area 2.

Simulation results for the scenarios above show that G-MPC can effectively utilize ancillary services from the buildings for frequency regulation purposes.

8 Conclusions

We addressed the problem of optimal design of an integrated energy managing system based on a supply-following strategy. In our framework, assume-guarantee contracts formalize the requirements of both the power grid and the commercial building subsystem, and specify their interface so as to allow for independent implementation of two model predictive control (MPC) schemes in a compositional fashion.

At the building level, we cast a robust optimal control problem to determine the baseline power consumption of the HVAC system and the amount of allowed power

consumption flexibility to maximize the building monetary incentive while satisfying its temperature and air mass flow requirements. At the grid level, we optimally control the ancillary service power flow within the buildings' flexibility, while integrating constraints such as ramping rates of ancillary service providers, maximum available ancillary power, and load forecast information. Simulation results show that commercial buildings can profitably provide ancillary services that can be effectively regulated at both the building and the grid levels by the proposed MPC scheme.

As a future work, we plan to further refine the communication interface between grid and buildings to incorporate a more realistic scenario, in which an electricity broker seeks rate offers from suppliers for "bundled" groups of customers and acts on their behalf. We also plan to develop a robust formulation of the grid MPC problem to address the uncertainties associated with imperfect load predictions.

Acknowledgements This work was supported in part by the TerraSwarm Research Center, one of six centers supported by the STARnet phase of the Focus Center Research Program (FCRP) a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

References

1. "Buildings energy data book." [Online]. Available: http://buildingsdatabook.eren.doe.gov/default.aspx
2. "U.S. Energy Information Administration, annual energy review," 2010. [Online]. Available: http://www.eia.gov/totalenergy/data/annual/consumption
3. (2012, Feb.) IBM ILOG CPLEX Optimizer. [Online]. Available: www.ibm.com/software/integration/optimization/cplex-optimizer/
4. A. Benveniste, B. Caillaud, A. Ferrari, L. Mangeruca, R. Passerone, and C. Sofronis, "Formal methods for components and objects," F. S. Boer, M. M. Bonsangue, S. Graf, and W.-P. Roever, Eds. Berlin, Heidelberg: Springer-Verlag, 2008, ch. Multiple Viewpoint Contract-Based Specification and Design, pp. 200–225.
5. A. Benveniste, B. Caillaud, D. Nickovic, R. Passerone, J.-B. Raclet, P. Reinkemeier, A. Sangiovanni-Vincentelli *et al.*, "Contracts for System Design," INRIA, Rapport de recherche RR-8147, Nov. 2012.
6. D. P. Bertsekas, "Nonlinear programming," 1999.
7. L. de Alfaro and T. A. Henzinger, "Interface automata." ACM Press, 2001, pp. 109–120.
8. A. Debs, *Modern power systems control and operation*. Kluwer Academic Publishers, Norwell, MA, 1988.
9. U. Helman, "Resource and transmission planning to achieve a 33% RPS in California–ISO modeling tools and planning framework," in *FERC Technical Conference on Planning Models and Software*, 2010.
10. A. Iannopollo, P. Nuzzo, S. Tripakis, and A. L. Sangiovanni-Vincentelli, "Library-based scalable refinement checking for contract-based design," in *Proc. Design, Automation and Test in Europe*, Mar. 2014.
11. B. Kirby, *Spinning reserve from responsive loads*. United States. Department of Energy, 2003.
12. S. Koch, J. Mathieu, and D. Callaway, "Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services," in *Proc. PSCC*, 2011, pp. 1–7.
13. P. Kundur, N. Balu, and M. Lauby, *Power system stability and control*. McGraw-hill New York, 1994, vol. 4, no. 2.

14. J. Lofberg, "Yalmip : A toolbox for modeling and optimization in MATLAB," in *Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004. [Online]. Available: http://users.isy.liu.se/johanl/yalmip
15. Y. Ma, F. Borrelli, B. Hencey, B. Coffey, S. Bengea, and P. Haves, "Model predictive control for the operation of building cooling systems," in *American Control Conference (ACC), 2010*. IEEE, 2010, pp. 5106–5111.
16. M. Maasoumy, A. Pinto, and A. Sangiovanni-Vincentelli, "Model-based hierarchical optimal control design for HVAC systems," in *Dynamic System Control Conference (DSCC), 2011*. ASME, 2011.
17. M. Maasoumy and A. Sangiovanni-Vincentelli, "Total and peak energy consumption minimization of building HVAC systems using model predictive control," *IEEE Design and Test of Computers*, Jul-Aug 2012.
18. M. Maasoumy, "Modeling and optimal control algorithm design for HVAC systems in energy efficient buildings," Master's thesis, EECS Department, University of California, Berkeley, Feb 2011. [Online]. Available: http://www.eecs.berkeley.edu/Pubs/TechRpts/2011/EECS-2011-12.html
19. M. Maasoumy, J. Ortiz, D. Culler, and A. Sangiovanni-Vincentelli, "Flexibility of commercial building HVAC fan as ancillary service for smart grid," in *IEEE Green Energy and Systems Conference*, Long Beach, USA, Nov. 2013.
20. M. Maasoumy, C. Rosenberg, A. Sangiovanni-Vincentelli, and D. Callaway, "Model predictive control approach to online computation of demand-side flexibility of commercial buildings HVAC systems for supply following," in *Proc. IEEE American Control Conf.*, Jun. 2014.
21. M. Maasoumy, B. M. Sanandaji, A. Sangiovanni-Vincentelli, and K. Poolla, "Model predictive control of regulation services from commercial buildings to the smart grid," in *Proc. IEEE American Control Conf.*, Jun. 2014.
22. M. Maasoumy Haghighi, "Controlling energy-efficient buildings in the context of smart grid: A cyber physical system approach," no. UCB, EECS-2013-244 - PhD Thesis, Dec 2013. [Online]. Available: http://www.eecs.berkeley.edu/Pubs/TechRpts/2013/EECS-2013-244.html
23. O. Maler and D. Nickovic, "Monitoring temporal properties of continuous signals," in *Formal Modeling and Analysis of Timed Systems*, 2004, pp. 152–166.
24. P. Nuzzo, J. B. Finn, A. Iannopollo, and A. L. Sangiovanni-Vincentelli, "Contract-based design of control protocols for safety-critical cyber-physical systems," in *Proc. Design, Automation and Test in Europe*, Mar. 2014.
25. P. Nuzzo and A. Sangiovanni-Vincentelli, "Let's get physical: Computer science meets systems," in *From Programs to Systems - The Systems Perspective in Computing Workshop, European Joint Conferences on Theory and Practice of Software (ETAPS)*, Apr. 2014.
26. P. Nuzzo, A. Sangiovanni-Vincentelli, X. Sun, and A. Puggelli, "Methodology for the design of analog integrated interfaces using contracts," vol. 12, no. 12, pp. 3329–3345, Dec. 2012.
27. P. Nuzzo, H. Xu, N. Ozay, J. Finn, A. Sangiovanni-Vincentelli, R. Murray, A. Donze, and S. Seshia, "A contract-based methodology for aircraft electric power system design," *IEEE Access*, vol. 2, pp. 1–25, 2014.
28. 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*. IEEE, 2010, pp. 5100–5105.
29. A. Pnueli, "The temporal logic of programs," in *18th Annual Symposium on Foundations of Computer Science*. IEEE Computer Society Press, 1977, pp. 46–57.
30. A. Sangiovanni-Vincentelli, "Quo vadis, SLD? Reasoning about the trends and challenges of system level design," *Proceedings of the IEEE*, vol. 95, no. 3, pp. 467–506, 2007.
31. A. Sangiovanni-Vincentelli, W. Damm, and R. Passerone, "Taming Dr. Frankenstein: Contract-Based Design for Cyber-Physical Systems," *European Journal of Control*, vol. 18, no. 3, pp. 217–238, Jun. 2012.
32. G. Strbac, "Demand side management: Benefits and challenges," *Energy Policy*, vol. 36, no. 12, pp. 4419–4426, 2008.

33. D. Todd, M. Caufield, B. Helms, M. Starke, B. Kirby, and J. Kueck, "Providing reliability services through demand response: A preliminary evaluation of the demand response capabilities of Alcoa Inc," *ORNL/TM*, vol. 233, 2008.
34. V. Vittal and A. Bergen, *Power systems analysis*. Prentice Hall, 1999.
35. A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical programming*, vol. 106, no. 1, pp. 25–57, 2006.
36. C. Woo, E. Kollman, R. Orans, S. Price, and B. Horii, "Now that California has AMI, what can the state do with it?" *Energy Policy*, vol. 36, no. 4, pp. 1366–1374, 2008.