Strengthening the Group: Aggregated Frequency Reserve Bidding with ADMM

Journal Article

Author(s):

Rey, Felix; Zhang, Xiaojing; Merkli, Sandro; Agliati, Valentina; [Kamgarpour, Maryam](#) Lygeros, John

Publication date:

2019

Permanent link:

<https://doi.org/10.3929/ethz-b-000268424>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

IEEE Transactions on Smart Grid, <https://doi.org/10.1109/TSG.2018.2841508>

Strengthening the Group: Aggregated Frequency Reserve Bidding with ADMM

Felix Rey, *Member, IEEE*, Xiaojing Zhang, *Member, IEEE*, Sandro Merkli, Valentina Agliati, Maryam Kamgarpour, *Member, IEEE*, and John Lygeros, *Fellow, IEEE*

Abstract—In a power grid, the electricity supply and demand must be balanced at all times to maintain the system’s frequency. In practice, the grid operator achieves this balance by procuring frequency reserves in an ahead-of-time market setting. During runtime, these reserves are then dispatched whenever there is an imbalance in the grid. Recently, there has been an increasing interest in engaging electricity consumers, such as plug-in electric vehicles or buildings, to offer such frequency reserves by exploiting their flexibility in power consumption. In this work, we focus on an aggregation of buildings that places a joint bid on a reserve market. The resulting shared decision is modeled as a large-scale optimization problem. Our main contribution is to show that the aggregation can make its decision in a computationally efficient and conceptually meaningful way, using the alternating direction method of multipliers (ADMM). The proposed approach exhibits several attractive features that include (i) the computational burden is distributed between the buildings; (ii) the setup naturally provides privacy and flexibility; (iii) the iterative algorithm can be stopped at any time, providing a feasible (though suboptimal) solution; and (iv) the algorithm provides the foundation for a reward distribution scheme that strengthens the group.

I. INTRODUCTION

COUNTRIES around the world are integrating renewable energy sources into their power systems. The electricity production of weather-dependent renewable energy sources, such as wind turbines and photovoltaics, is volatile. This increases the uncertainty in the grid operation and challenges the balance of electricity supply and demand. For maintaining the active power balance and regulating the grid frequency, grid operators procure ancillary services, such as reserved generation capacities [1]. Traditionally, these reserves are provided by conventional generators. However, the rapid integration of renewable energy sources calls for an increased balancing effort [2]. Recently, it has been shown that if properly aggregated and controlled, flexible loads can be engaged to provide frequency regulation at low cost and small environmental impact [3]. Consumers with sizeable thermal storage, such as buildings, are particularly suited to contribute to the grid balance as their electricity consumption can be shifted in time with negligible impact on occupant comfort [4], [5].

In many regions, consumers provide frequency reserves by offering them to the grid operator in a market setting. In

Switzerland, the offered reserve size needs to be constant over the bidding period and exceed a certain minimum amount [6]. These requirements reduce the coordination effort for the grid operator by simplifying the scheduling and avoiding a large number of small bidders. The minimum bid typically ranges from 1 MW to 10 MW, e.g., it is 5 MW in Switzerland [6, Section 2.3.1], which is too large for a single building [7]. To still participate in the reserve market, buildings can cooperate [8], which requires solving a large-scale optimization problem for obtaining a time-constant joint bid. A centralized solution of such a problem faces several issues: (i) a shared investment is necessary to establish and maintain a central computation unit; (ii) solving the resulting large-scale optimization problem can be computationally challenging, which effectively limits the aggregation size; (iii) the building operators need to share sensitive information, such as details on the construction material and occupants’ behavior; and (iv) each local configuration change needs to be communicated to the central unit.

In this paper, we address these issues by using the alternating direction method of multipliers (ADMM) [9], [10] to distribute the original large-scale optimization problem among the individual buildings. Each building then solves a localized problem and participates in a global mediation process until a consensus solution is reached. The distribution of the computational burden provides natural parallelism and therefore makes it possible to use a simple central processing unit that only handles coordination tasks. Further, we present a modified communication scheme that avoids a central unit altogether. We also show that increased privacy and flexibility follows from the distribution of information, and we demonstrate that we obtain a joint bid satisfying all constraints even if we terminate ADMM prematurely. Finally, given that a reserve bid is accepted on the market, we show a reward distribution scheme that reduces provision imbalance within the aggregation, strengthening the group’s capability for providing reserves.

Related Literature

Frequency reserve provision by buildings receives considerable attention, both from the power systems community [3], [8], [11]–[17] and the control community [7], [18]–[26]. In [8], [12], the primary goal is to determine the optimal reserve capacity that an aggregation of buildings can provide while ensuring occupant comfort. In [13], [14], frequency reserves are provided by modulating the fan speed of air handling units in commercial buildings. However, rather than giving a priori guarantees on comfort constraints as in [8], [12], the authors in [13], [14] test constraint satisfaction a posteriori

F. Rey `rey*`, S. Merkli `smerkli*`, M. Kamgarpour `mkamgar*` and J. Lygeros `lygeros*` are with the Automatic Control Laboratory at ETH Zurich, Switzerland, `*@control.ee.ethz.ch`. Our research is supported by the European Research Council under ERC Starting Grant CONENE and by ABB Corporate Research under grant 2017-1224/01. X. Zhang `xiaojing.zhang@berkeley.edu` is with the MPC Lab at UC Berkeley, USA. V. Agliati graduated from Politecnico di Milano in 2017 and contributed to this work during an exchange at ETH Zurich.

through experimental trials. Our work follows [8], [12], using the mathematical tools for robust constraint satisfaction from [18], [22]. In contrast to [8], [12], our leading objective is a distributed and decentralized solution of the multi-building problem. As opposed to [17], [27], where voltage regulation affects the reactive power balance, we concentrate on frequency regulation and active power. Aside from our focus on market bidding, [16], [24]–[26] analyze the real-time balance in the power grid. In [25], [26], the grid is modeled explicitly, exceeding our work as we purely rely on the reserve demand as requested by the grid operator. Incorporating such grid dynamics can be worthwhile, particularly for spatially separated aggregations. While [16] guides the demand response based on time-varying energy prices, [24] considers a setup where previously agreed reserves are utilized. Hence, [24] augments our work particularly well, as we focus on a market where such reserve agreements are made.

In [11], [28]–[30], ADMM is studied in the context of demand response. In [28], a distributed procedure coordinates different temperature zones inside a building. In [29], ADMM exploits the flexibility of several consumers for shaping their combined power intake. In [11], heterogeneous agents are aggregated. Finally, in [30], small thermostatically controlled loads are orchestrated to participate in real-time energy markets. We differ from these publications by focusing on ahead-of-time reserve bidding while the actual reserve request is still unknown. We further set ourselves apart by developing additional tools that improve the conceptual utility of our approach, such as a fully-decentralized computation scheme, the permanent availability of a feasible solution, and a profit allocation that places an incentive to strengthen the group.

Notation

For $x, y \in \mathbb{R}^n$, we use $(x; y) = [x^\top y^\top]^\top \in \mathbb{R}^{2n}$. Inequalities between vectors, e.g., $x \leq y$, as well as inclusions, e.g., $x \in [-y, y]$, are understood element-wise. We write $\{x_i\}_i$ for the tuple $(x_1, x_2, \dots)$, where the context determines the range of i . We denote the Kronecker product as $\otimes$. We use the identity matrix I_n of dimension $n \times n$. $0_{n \times m}$ or $1_{n \times m}$ are $n \times m$ matrices with all elements 0 or 1, respectively; we omit the subscripts if the dimension is clear from context.

II. FREQUENCY RESERVE PROVISION

We first introduce the model of a single building and describe the participation in the reserve market. We then show the joint bidding procedure for aggregations of buildings.

A. Building Model and Reserve Market

Following [15], [31], [32], we consider the building model

$$x^{k+1} = Ax^k + B(u^k + \Delta u^k) + Ev^k, \quad (1)$$

where we use the admissible state $x^k \in \mathcal{X} \subseteq \mathbb{R}^n$, the admissible input $(u^k + \Delta u^k) \in \mathcal{U} \subseteq \mathbb{R}^m$, and the external disturbance $v^k \in \mathbb{R}^q$. We use the discrete time $k = 1, 2, \dots, N$ with horizon N . The state x^k models temperatures in rooms, walls, floors, and ceilings across the building. We use $\mathcal{X}$ to model convex comfort constraints, e.g., to limit room temperatures

between 21°C and 25°C . The system matrix A describes the temperature diffusion with a thermal resistance-capacitance model as in [31, Sec. 4.2]. The actuation (heating, cooling, and ventilation) is described by a nominal input u^k and a variable term Δu^k . The variable term is used later for adjusting the buildings' power consumption. The convex set $\mathcal{U}$ models actuation limits and saturation. Finally, we use Ev^k to describe the effect of external disturbances (occupancy, solar radiation, and ambient temperature). Section IV shows a simulation study with detailed modeling examples. The mechanisms behind building-internal temperature regulation are specifically addressed in [31], [32]. As we focus on the collaboration of different buildings in this paper, we suffice with using the abstraction (1). Indeed, any linear system A, B, E with convex constraints $\mathcal{X}, \mathcal{U}$ is compatible with the subsequent analysis. We also simplify the setup by assuming perfect prediction for v^k , noting that uncertainty can easily be integrated into our approach by using the robust control methods in [33]–[35]. With $x := (x^1; \dots; x^{N+1})$, $u := (u^1; \dots; u^N)$, and similar definitions for $\Delta u, v$ (ranging from 1 to N), we rewrite (1) as

$$x = Ax^1 + B(u + \Delta u) + Ev \quad (2)$$

with $x \in \mathcal{X}$ and $(u + \Delta u) \in \mathcal{U}$. We assemble $A, B, E, \mathcal{X}, \mathcal{U}$ from $A, B, E, \mathcal{X}, \mathcal{U}$ as shown in [36].

In Switzerland, similar to most of Europe, reserve capacities are procured in a market where reserve providers bid their capacity in weekly or daily blocks. We denote the time-varying reserve bid as $y = (y^1; \dots; y^N)$, which describes the maximum-possible change in active power consumption, measured in kilowatt. We focus on a day-ahead market, and we consider a symmetric reserve range $[-y, y]$, meaning that the building needs to shift its consumption equally in both directions. The limitation to symmetric provision follows the current policy in Switzerland [6]; the extension to asymmetric reserves is straightforward. We further require a constant bid over the bidding period, i.e., $y^1 = y^2 = \dots = y^N$, which is also standard in several countries. The restriction to time-constant reserve simplifies the scheduling for the grid operator and prevents an imbalance over the bidding interval. Once a bid is accepted, the building receives a reward of the form $p^\top y$, where $p \in \mathbb{R}^N$ is a price parameter that we assume to be known, e.g., from previous market clearings. During actual runtime, the grid operator then requests a reserve demand $s = (s^1; \dots; s^N) \in [-y, y]$, and the building is obliged to adjust its power consumption accordingly. The challenge in ahead-of-time bidding is choosing y while s is still unknown. Fig. 1 summarizes the involved variables.

Fig. 1. Decision variables $x, u, \Delta u, y$, which can be determined by the building, and external variables s, v , which cannot be influenced.

After a bid is accepted at the reserve market, the building follows the request s by appropriately changing its input from u to $u + \Delta u$. More precisely, following [31], [34], [37], we use the conversion factor $\eta \in \mathbb{R}^m$ for translating input quantities (e.g., heating levels) into power consumption (measured in kilowatt). We choose the variable input Δu such that $s^k = \eta^\top \Delta u^k$ for all k . Fig. 2 illustrates the separation between reserve bidding and the actual provision procedure.

Fig. 2. Bidding process and reserve provision. *Left:* Ahead of time, the building offers a bid y on the reserve market. If the bid gets accepted, it receives the reward $p^\top y$. The actual reserve request s^k remains unknown at this point. *Right:* During runtime, the building follows the reserve request s^k within the previously offered range $[-y^k, y^k]$ by adjusting its variable input Δu^k .

Our approach focuses on the bidding process, particularly on determining a satisfiable and financially profitable bid for the reserve market. Additionally, our bidding method results in a rudimentary procedure to follow the reserve request during runtime, meaning that we relate Δu to s , which is further refined in the following section. Beyond the scope of our paper, our provision procedure can be complemented with more sophisticated and iterative reserve tracking techniques as in [24]–[26], which can exploit newly-available information or incorporate grid models to fine-tune the supply.

B. Individual Bidding Problem

Typically, buildings need to increase their nominal consumption u for being able to react to symmetric reserve requests [8]. The building's objective is to optimally balance the provision reward $p^\top y$ and the nominal electricity cost $c^\top u$, i.e., it minimizes its overall cost $c^\top u - p^\top y$. Here, the vector c is composed of subvectors $c_k = \tilde{c}_k \eta$, where $\tilde{c}_k \in \mathbb{R}$ is the electricity price, and η , as introduced before, converts input signals to power consumption. Since the reserve request $s \in \mathbb{R}^N$ is unknown at the time of bidding, we introduce recourse into the formulation by choosing u and Δu to be closed-loop policies

$$u(\cdot) : \mathbb{R}^N \rightarrow \mathbb{R}^{Nm} \quad s \mapsto u \quad (3a)$$

$$\Delta u(\cdot) : \mathbb{R}^N \rightarrow \mathbb{R}^{Nm} \quad s \mapsto \Delta u, \quad (3b)$$

which are parametrized functions that determine the actual input only when s becomes known during runtime [36], [38], [39]. In other words, instead of determining the static values $u, \Delta u$ that prepare for the worst-case request s , we try to find the optimal functional relation of $u, \Delta u$ to s . Compared to the conventional strategy, closed-loop policies offer a larger degree of freedom and therefore potentially less conservative reserve bids [20]. We choose the policies to

be causal such that $u^k, \Delta u^k$ only depend on $s^1, \dots, s^k$. The building-individual reserve provision problem is

$$\min_{u(\cdot), y} \quad c^\top u(0) - p^\top y \quad (4a)$$

$$\text{s.t.} \quad (u(\cdot), y) \in \mathcal{C} \quad (4b)$$

$$y^1 = y^2 = \dots = y^N. \quad (4c)$$

For simplicity, we use the nominal electricity cost $c^\top u(0)$, which is based on the case $s = 0$ where the grid operator does not request any reserves. Due to $\Delta u(0) = 0$, the variable input is not part of the objective. Alternatively, it is possible to use the expected cost $\mathbb{E}[c^\top (u(s) + \Delta u(s))]$, provided we have a probabilistic description of the demanded reserve s [20]. Further, we choose not to model the minimum reserve constraint $y^k \geq y_{\min}$. Having such a constraint may lead to overall increasing costs compared to not providing reserves at all. Instead, we purely maximize financial profit, and we assume that the resulting bid is only placed at the market if it exceeds the minimum size. Omitting the minimum reserve constraint also benefits the construction of the distributed setup, as we will clarify in Remark 2. The set $\mathcal{C}$ in (4b) captures the building dynamics and the reserve provisioning mechanism as introduced before. It encapsulates building-internal information, e.g., the state x , which in consequence of (3) also becomes a policy. We define $\mathcal{C}$ as

$$\mathcal{C} = \left\{ (u(\cdot), y) \in (\mathcal{F} \times \mathbb{R}^N) \mid \exists (\Delta u(\cdot), x(\cdot)) \in \mathcal{F} \times \mathcal{F} \text{ s.t.} \right. \\ \left. \begin{aligned} x(s) &= \mathbf{A}x^1 + \mathbf{B}[u(s) + \Delta u(s)] + \mathbf{E}v \\ (x(s), u(s) + \Delta u(s)) &\in \mathcal{X} \times \mathcal{U} \\ s &= (\mathbf{I}_N \otimes \eta^\top) \Delta u(s) \end{aligned} \right\} \forall s \in [-y, y] \}, \quad (5)$$

where $\mathcal{F}$, in the most general case, is the infinite-dimensional space of causal functions of type (3). Due to the optimization over policies and the infinite number of constraints, solving (4) is intractable in general. However, if we restrict $\mathcal{F}$ to contain linear functions only, then we can reformulate (4) as a convex optimization problem with a finite number of constraints; see Appendix A and [22, Section 4] for an example. Given such an appropriate choice of $\mathcal{F}$, the constraint set $\mathcal{C}$ and therefore the individual bidding problem (4) is finite-dimensional and convex. Indeed, finite dimensionality and convexity of $\mathcal{C}$ are the only requirements for the analysis below. Hence, our approach is compatible with different policy spaces and building models. Further, it can be used for plug-in electric vehicles, house-mounted batteries, and any energy storage device that has linear dynamics $\mathbf{A}, \mathbf{B}, \mathbf{E}$ and convex constraints $\mathcal{X}, \mathcal{U}$.

Remark 1. In some markets, the price p is also a decision variable, which increases the computational cost significantly. For keeping the computation simple, we can solve (4) for different values of p , resulting in a price-reserve curve that is bid at the market. This strategy is known as conditional bidding.

C. Aggregated Bidding Problem

On their own, buildings are typically not able to meet the required minimum bid size [15]. For overcoming this, buildings can collaborate in an aggregation. We consider M buildings, indexed by $b = 1, \dots, M$, and we denote the reserve bids as y_b , where each building's bid has the form $y_b = (y_b^1; \dots; y_b^N)$.

Similarly, we use u_b , c_b , and C_b for the individual input, cost function, and constraint set. Buildings are compatible as long as they agree on the bidding horizon N , which means that they can use different models in C_b . We use non-calligraphic uppercase symbols for quantities associated with the entire aggregation, e.g., we define the aggregated reserve $Y := \sum_b y_b$, where $Y = (Y^1; \dots; Y^N)$. The aggregated bidding problem is

$$\min_{Y, \{u_b(\cdot), y_b\}_b} \sum_{b=1}^M c_b^\top u_b(0) - p^\top Y \quad (6a)$$

$$\text{s.t.} \quad (u_b(\cdot), y_b) \in C_b \quad \forall b \quad (6b)$$

$$Y = \sum_{b=1}^M y_b \quad (6c)$$

$$Y^1 = Y^2 = \dots = Y^N. \quad (6d)$$

In the aggregated problem, the time-constant constraint (6d) applies to the aggregated bid Y , which gives each building the freedom of time-varying bids y_b as long as their sum remains constant. This additional freedom increases the joint bid, as compared to the trivial strategy of pooling individual bids obtained with (4). The potential increase in Y comes at the cost of requiring a deeper collaboration between the buildings, as each building has to shape its offer y_b while considering the bids of all other buildings. If the time-constant constraints (4c) and (6d) would not be present, then (6) would simply decompose into separate instances of (4), i.e., there would be no need for collaboration. When the buildings agree on a joint bid Y that is also accepted at the reserve market, then the grid operator can request his demand s from the aggregation during runtime. Each building then changes its consumption proportionally to its share y_b of the original offer Y . Beyond the scope of our paper, it is possible to successively reallocate the provision responsibilities within the aggregation during runtime, which can be especially useful when taking a grid model and therefore local imbalances into account [25], [26].

A direct solution of (6) has several downsides: (i) Due to the large dimensionality, we require a powerful central processing unit, which demands financial resources and restricts the aggregation size through limited computational capabilities. (ii) The buildings need to disclose C_b , potentially including sensitive information, which also implies that they cannot change C_b without informing the central unit. (iii) The buildings delegate the decision process, i.e., they have to rely on the central unit for finding a good solution to (6).

III. AGGREGATED BIDDING WITH ADMM

We use ADMM to decompose (6) into a sequence of smaller problems, resulting in conceptual benefits over a direct solution approach. We rely on ADMM due to its mild convergence requirements, e.g., we do not need strong convexity.

A. General ADMM Formulation

ADMM is a first-order optimization method that solves problems where the decision variables are partitioned into two groups. Typically, these problems have the form

$$\min_{z,w} \{ f(z) + g(w) \mid \text{s.t. } \mathcal{A}(z) + \mathcal{B}(w) = 0 \} \quad (7)$$

with real-valued functions f, g and affine operators $\mathcal{A}, \mathcal{B}$.

As shown in Alg. 1, ADMM solves (7) by alternatingly searching for a saddle point of the augmented Lagrangian

$$\mathcal{L}_\rho(z, w, \lambda) := f(z) + g(w) + \lambda^\top \epsilon + \frac{\rho}{2} \|\epsilon\|_2^2, \quad (8)$$

where λ is the Lagrange multiplier, $\epsilon := \mathcal{A}(z) + \mathcal{B}(w)$ is a residual, and $\rho > 0$ is a user-chosen penalty parameter. For any initialization of (z, w, λ) , and any positive value of ρ , ADMM converges to a solution of (7), given that a solution exists, and f, g have a nonempty, closed, and convex epigraph [10].

Algorithm 1 ADMM

repeat 1: $z \leftarrow \arg \min_z \mathcal{L}_\rho(z, w, \lambda)$

2: $w \leftarrow \arg \min_w \mathcal{L}_\rho(z, w, \lambda)$

3: $\lambda \leftarrow \lambda + \rho \frac{\partial}{\partial \lambda} \mathcal{L}_\rho(z, w, \lambda)$

until satisfaction of a stopping criterion

B. Aggregated Bidding with ADMM

For writing (6) in the form of (7), we first need to relate the decision variables $Y, \{u_b(\cdot), y_b\}_b$ to the two-component splitting z, w . Further, we need to associate each constraint in (6) to one of the two groups, and we need to decide on a coupling relation $\mathcal{A}, \mathcal{B}$. Our goal is to associate w with the aggregation-specific variable Y , such that z only contains building-individual variables $\{u_b(\cdot), y_b\}_b$, making the z -update in Alg. 1 decompose among buildings. When pursuing such a splitting, we struggle with assigning (6c) to one side, as it connects the aggregation-specific and building-individual variables. We overcome this difficulty by making (6c) depend on the auxiliary decision variable $\{\bar{y}_b\}_b$, which we constrain to be equal to $\{y_b\}_b$. Later on, it will become clear that having these extra variable does not increase the computational burden. With $\{\bar{y}_b\}_b$, we then can separate between the aggregation-specific variables $w := (Y, \{\bar{y}_b\}_b)$, and the building-individual variables $z := \{u_b(\cdot), y_b\}_b$, using $\{\bar{y}_b\}_b = \{y_b\}_b$ as coupling relation. The resulting problem partition, shown in Fig. 3, is an equivalent reformulation of (6).

$$\begin{array}{ll} \min_{\substack{\{u_b(\cdot), y_b\}_b, \\ Y, \{\bar{y}_b\}_b}} & \boxed{\sum_b c_b^\top u_b(0)} \quad - p^\top Y \quad \text{building-individual} \\ \text{s.t.} & \boxed{(u_b(\cdot), y_b) \in C_b \quad \forall b} \quad \text{aggregation-specific} \\ & \boxed{Y = \sum_b \bar{y}_b} \\ & \boxed{Y^1 = Y^2 = \dots = Y^N} \\ & \boxed{\bar{y}_b - y_b = 0 \quad \forall b} \end{array}$$

Fig. 3. Separation of (6) into building-individual and aggregation-specific terms, using the auxiliary variable $\bar{y}_b = y_b$ for $b = 1, \dots, M$.

By comparing the partitioned problem in Fig. 3 to the general ADMM problem (7), we obtain the definitions

$$f(z) := \sum_b (c_b^\top u_b(0) + \mathcal{I}_{C_b}(u_b(\cdot), y_b)) \quad (9a)$$

$$g(w) := \mathcal{I}_{Y=\sum_b \bar{y}_b}(Y, \{\bar{y}_b\}_b) + \mathcal{I}_{Y^1=\dots=Y^N}(Y) - p^\top Y \quad (9b)$$

$$\mathcal{A}(z) + \mathcal{B}(w) = 0 \Leftrightarrow \{\bar{y}_b\}_b = \{y_b\}_b, \quad (9c)$$

where indicator functions are denoted as $\mathcal{I}(\cdot)$. Inserted into (8) and Alg. 1, these definitions lead to our proposed ADMM formulation, which we show in Alg. 2.

Algorithm 2 ADMM for Reserve Provision**repeat**

① building-individual reserve proposition ($b = 1, \dots, M$)
 $(u_b(\cdot), y_b) \leftarrow \arg \min_{u_b(\cdot), y_b} c_b^\top u_b(0) - \lambda_b^\top y_b + \frac{\rho}{2} \|\bar{y}_b - y_b\|_2^2$
s.t. $(u_b(\cdot), y_b) \in \mathcal{C}_b$

② aggregation step
 $(Y, \{\bar{y}_b\}_b) \leftarrow \arg \min_{Y, \{\bar{y}_b\}_b} \sum_{b=1}^M (\lambda_b^\top \bar{y}_b + \frac{\rho}{2} \|\bar{y}_b - y_b\|_2^2) - p^\top Y$
s.t. $Y = \sum_b \bar{y}_b$
 $Y^1 = Y^2 = \dots = Y^N$

③ Lagrangian mediation ($b = 1, \dots, M$)
 $\lambda_b \leftarrow \lambda_b + \rho(\bar{y}_b - y_b)$

until satisfaction of the stopping criterion

Alg. 2 alternates between a building-individual step ①, an aggregation-specific step ②, and Lagrangian mediation ③. For the initialization, we can set Y , $\bar{y}_b$, and λ_b to zero. Prop. 1, proven in Appendix B-A, constitutes the use of Alg. 2.

Proposition 1. *If $\mathcal{C}_b$ is nonempty, finite-dimensional, and convex for each building b , then Alg. 2 converges to an optimal solution of the aggregated bidding problem (6).*

We interpret Alg. 2 as a negotiation process between the buildings and the aggregator. In ①, each building proposes a reserve capacity y_b , maximizing its profit. The aggregator, which receives these propositions, responds with the desired reserve amount $\bar{y}_b$, chosen in ② for achieving a time-constant total reserve Y . The Lagrangian update ③ forms a price mechanism that mediates between proposition y_b and demand $\bar{y}_b$ to reach a consensus. Fig. 4 illustrates the information exchange.

Fig. 4. Information exchange between the buildings and the aggregator within one iteration of Alg. 2. Illustrated is the case of $M = 2$.

The ADMM-based approach shows conceptual advantages when compared to a centralized solution of (6). Due to the building-individual execution of ①, the constraint sets $\mathcal{C}_b$ remain private, and the buildings can change them self-organized. Also, the execution parallelizes naturally. Other benefits of Alg. 2 derive from the simplicity of its aggregation step, which is described in the following proposition.

Proposition 2. *The aggregation step ② has the solution*

$$\Omega := \frac{1}{M} \sum_{b=1}^M (\rho y_b - \lambda_b) \quad (10a)$$

$$Y = \frac{M}{\rho N} \mathbf{1}_{N \times N} (\Omega + p) \quad (10b)$$

$$\bar{y}_b = \frac{1}{\rho} (\rho y_b - \lambda_b - \Omega) + \frac{1}{M} Y \quad \forall b, \quad (10c)$$

where Ω is called the aggregate.

Prop. 2, which we prove in Appendix B-B, makes it possible to replace the aggregation step with a set of closed-form algebraic equations. This signifies that the computational complexity of ② is negligible compared to the first (parallel)

algorithm step. As a consequence, the increased computational burden for large aggregations is fully absorbed by the equally growing level of parallelism. The ADMM-based framework also lowers the barriers to integrate new buildings into an existing aggregation. The reason is the similarity between the individual problem (4) and the building-individual step ① in Alg. 2, i.e., new buildings that switch from individual bidding practically keep performing the same optimization as before.

Remark 2. *If we model a minimum bid constraint $y^k \geq y_{\min}$ in (4), and accordingly $Y^k \geq y_{\min}$ in (6), then the aggregation step contains inequality constraints and therefore does not have a closed-form solution anymore.*

C. Decentralized Computation

Alg. 2 is flexible against configuration changes, avoids concentration of private information, and behaves well from a computational perspective. Still, the algorithm requires the acquisition and maintenance of a central computation facility, which poses an organizational burden. For overcoming this, we recognize that in the aggregation step (10), only the aggregate Ω requires aggregation-global information. More specifically, we notice that when Ω is available, then each building can evaluate the previously centralized operations (10b) and (10c) individually. Towards utilizing this, we define the partial aggregate

$$\Omega_i := \frac{1}{M} \sum_{b=1}^i (\rho y_b - \lambda_b), \quad (11)$$

which we can accumulate by serially handing it from building to building. Each building b then adds its own contribution $\frac{1}{M}(\rho y_b - \lambda_b)$, assuming that the aggregation size M is known. Once we have reached $\Omega = \Omega_M$, the buildings circulate the result back for making it known to their neighbors. Fig. 5 illustrates the communication procedure for the case of four buildings. In general, we can use any connected communication graph that covers the aggregation.

Fig. 5. Decentralized information exchange to make the aggregate Ω available to all buildings. Illustrated is the case of $M = 4$.

Alg. 3 summarizes the resulting decentralized formulation. The algorithm is obtained by applying the previously described decentralization procedure to Alg. 2, and by rotating its steps from ①-②-③ to ②-③-①, resulting in a grouped building-individual step ① and a serial communication step ②. Alg. 3 is an equivalent reformulation of Alg. 2, i.e., both algorithms produce the same sequence of iterates when initialized the same. Therefore, all properties that are discussed in this paper equally apply to both algorithms. In subsequent sections, we again focus on Alg. 2, utilizing its more explicit aggregator-building composition to simplify the arguments that follow.

Algorithm 3 Decentralized ADMM for Reserve Provision**repeat**① building-individual updates ($b = 1, \dots, M$)

$$\bar{y}_b \leftarrow \frac{1}{\rho}(\rho y_b - \lambda_b - \Omega) + \frac{1}{\rho N} \mathbf{1}_{N \times N} (\Omega + p)$$

$$\lambda_b \leftarrow \lambda_b + \rho(\bar{y}_b - y_b)$$

$$(u_b(\cdot), y_b) \leftarrow \arg \min_{u_b(\cdot), y_b} c_b^\top u_b(0) - \lambda_b^\top y_b + \frac{\rho}{2} \|\bar{y}_b - y_b\|_2^2$$

$$\text{s.t. } (u_b(\cdot), y_b) \in \mathcal{C}_b$$

② serial communication

accumulate Ω_i and circulate Ω as shown in Fig. 5**until** satisfaction of the stopping criterion**end** obtain the aggregated bid Y with (10b)**D. Feasible Extraction**

In Alg. 2, the partial results of ① and ② satisfy their respective constraints; however are not compatible as it is $y_b \neq \bar{y}_b$ before convergence. In other words, ADMM only finds a solution that jointly satisfies all constraints in the limit, potentially requiring a large number of iterations [40]. Here, we overcome this difficulty with a method for extracting a jointly feasible solution after any number of ADMM iterations. This extraction procedure has great value in practice as it allows for terminating the algorithm anytime, e.g., when we reach a given time limit. The following proposition describes the procedure.

Proposition 3. Given $\{u_b(\cdot), y_b\}_b$ from any iteration of Alg. 2, we obtain a suboptimal solution $(Y^F, \{u_b^F(\cdot), y_b^F\}_b)$ that satisfies all constraints in (6) with

$$Y^F := \mathbf{1}_{N \times 1} \min_k \left(\sum_{b=1}^M y_b^k \right) \quad (12a)$$

$$(y_b^k)^F := \frac{(Y^k)^F}{\sum_{j=1}^M y_j^k} y_b^k \quad \forall b, k \quad (12b)$$

$$u_b^F(\cdot) := \arg \min_{u_b(\cdot)} \{c_b^\top u_b(0) \text{ s.t. } (u_b(\cdot), y_b^F) \in \mathcal{C}_b\} \quad \forall b. \quad (12c)$$

We show the proof of Prop. 3 in Appendix B-C. According to (12a), we obtain the feasible combined bid Y^F by clipping the accumulated reserve offer to its lowest value over time. With (12b), the individual reserve contributions y_b are scaled down to y_b^F , which then satisfy $Y^F = \sum_b y_b^F$. In (12c), we reoptimize the input. Due to $y_b^F \leq y_b$ the previous policy $u_b(\cdot)$ also remains feasible; however is less efficient. The more algorithm iterations we perform before extracting a solution, the lower the overall cost $\sum_b c_b^\top u_b^F(0) - p^\top Y^F$. In Section IV-B, we empirically analyze how many iterations are necessary to come reasonably close to the optimum.

E. Reward Distribution

When the grid operator accepts a reserve bid Y , the aggregation receives the reward $R := p^\top Y$. We can distribute this reward in proportion to the offered reserves, i.e., building b gets

$$r_b := p^\top y_b. \quad (13)$$

It is $R = \sum_b r_b$, signifying that we do not distribute more or less than the total reward. A drawback of this approach is that

it does not account for the relative importance of a reserve contribution, which we clarify at the example in Fig. 6.

Fig. 6. Illustrative example where seven buildings place an imbalanced reserve bid over a period of $N = 12$ hours. During most of the time, six similar buildings (black) jointly offer reserves $y_1 = y_2 = \dots = y_6$, while at $k = k_{crit}$ a single building (red) determines the total reserve $Y^k = y_7^k$.

In Fig. 6, a single building determines the combined reserve offer at the critical time $k = k_{crit}$. We construct the example by using dynamic-free buildings and by constraining the black buildings' provision capability to zero at k_{crit} . Due to the time-constant constraint, the aggregation cannot provide any reserves if the red building leaves the group. Conversely, if a black building exits the aggregation, the reserve offer decreases only marginally (if at all), since the remaining black buildings can compensate for it. Hence, due to its provision at k_{crit} , the red building is more important. With the rewarding scheme (13) however, the importance of provision during a particular time instance is not taken into account. The red building even earns less than any of the black buildings, e.g., for a time-constant price p it only gets $6/11$ of each black building's profit. We propose a Lagrangian-based approach to devise a reward allocation mechanism that better reflects the contribution of each building to the performance of the aggregation. As noted in Section III-B, the Lagrange multipliers $\lambda_b = (\lambda_b^1; \dots; \lambda_b^N)$ depend on time k , building b , and they also change with each ADMM iteration. Within Alg. 2, they orchestrate the buildings for achieving a constant combined reserve offer. For a particular time instance, the intuition is that when the aggregation demand continually exceeds a building's proposition, i.e., $\bar{y}_b^k - y_b^k > 0$, then the Lagrangian update

$$\lambda_b^k \leftarrow \lambda_b^k + \rho(\bar{y}_b^k - y_b^k)$$

accumulates λ_b^k towards a large value. The more difficult the provision, the larger λ_b^k grows before we reach a consensus. Fig. 7 illustrates the mechanism for $\lambda_b^k \geq 0$. In case of negative Lagrange multipliers, e.g., due to the initialization, it can be seen from inspecting ① and ② in Alg. 2 that we recover a non-negative Lagrange multiplier quickly.

Following Fig. 7, the Lagrange multipliers indicate the time instances where the reserve provision is difficult. We use this effect for increasing the reward during such critical hours. Specifically, we define the Lagrangian-based reward allocation

$$r_b^\Lambda := \Lambda^\top y_b, \quad (14)$$

where $\Lambda \in \mathbb{R}^N$ is an alternative price vector that replaces the static reward p in (13). We describe the relation between Λ and the Lagrange multipliers in the following proposition.

Fig. 7. Lagrangian mediation for $\lambda_b^k \geq 0$. If $\bar{y}_b^k > y_b^k$, then λ_b^k increases by $\Delta\lambda_b^k$ between ADMM iterations. This increase produces an incentive for raising y_b^k and lowering $\bar{y}_b^k$, which reduces the gap between them. The amount of change in y_b^k and $\bar{y}_b^k$ depends on the remaining constraints and objectives.

Proposition 4. For Alg. 2, the following is true.

- (i) After each algorithm iteration, all buildings have the same Lagrange multiplier, i.e., $\lambda_1 = \dots = \lambda_M$. We write $\Lambda := \lambda_b$ for any b , where $\Lambda = (\Lambda^1; \dots; \Lambda^N)$.
- (ii) The Lagrangian update (step ③ in Alg. 2) can be written as $\Lambda \leftarrow \frac{\rho}{M} Y - \Omega$.
- (iii) For a converged solution $(\Lambda, Y, \{y_b\}_b)$, the Lagrangian-based allocation $r_b^\Lambda = \Lambda^\top y_b$ satisfies $\sum_b r_b^\Lambda = R$.
- (iv) For a feasible suboptimal solution $(Y^F, \{y_b^F\}_b)$ that is obtained with Prop. 3, and

$$\Omega^F := \frac{\rho}{M} Y^F - \Lambda \quad (15a)$$

$$\Lambda^F := \frac{1}{N} 1_{N \times N} (\Omega^F + p) - \Omega^F, \quad (15b)$$

where Λ is obtained from the last algorithm iteration, the reward allocation $r_b^\Lambda = (\Lambda^F)^\top y_b^F$ satisfies $\sum_b r_b^\Lambda = R$.

We show the proof of Prop. 4 in Appendix B-D. In (i), we define the alternative price vector Λ as the building-invariant Lagrange multiplier. While the Lagrange multiplier is the same for all buildings, it still varies over the bidding time, indicating hours where provision is difficult. In Statement (ii), we provide a simple equation for Λ . Statement (iii) verifies that the allocation scheme distributes not more or less than the total reward R , and Statement (iv) makes it possible to use Lagrangian-based allocation after feasible extraction. We interpret the building-invariance of the Lagrange multipliers as follows: λ_b mediates between building proposal y_b and aggregator demand $\bar{y}_b$. However, the aggregator does not use any building-private information. In other words, it decides on the reserve demand $\bar{y}_b$ without being able to differentiate between the buildings. Therefore, it always asks for the same change $\bar{y}_b - y_b$, resulting in the same multipliers.

The Lagrangian-based allocation provides an incentive for reducing provision imbalance. Hence, it strengthens the group's capability for providing reserves, meaning that the aggregation grows more resilient against failures and configuration changes. However, in some situations, a pure utilization of the new allocation scheme can be overly extreme for being economically acceptable, which becomes clear from Fig. 8.

Fig. 8. Lagrange multiplier Λ , resulting from the example in Fig. 6. The reward concentrates in one point since a constraint limits the reserve provision only in a single time instant and the buildings are dynamic-free.

While the extreme reward concentration in Fig. 8 is specific to the constructed example, unrewarded hours can also occur in more generic setups. For alleviating this, we can use a mixed reward $\alpha r_b + (1 - \alpha) r_b^\Lambda$ with an acceptable weight $\alpha \in [0, 1]$.

IV. SIMULATION STUDY

We apply ADMM-based reserve bidding in simulation, considering a 24-hour bidding period with hourly discretization. We use a test set of 300 building models, each generated by random perturbation from one of the following prototypes.

- (i) Small building, obtained from [33, Section 4.5]
 - 3 states for temperatures, using comfort constraints between 21°C and 25°C when the building is occupied
 - 4 constrained inputs for radiator, cooled ceiling, floor heating, and mechanical ventilation
 - 3 modeled disturbances for outside temperature, solar radiation, and occupancy
- (ii) Medium-sized building, obtained from [37]
 - 33 states for temperatures, using comfort constraints between 20°C and 28°C when the building is occupied
 - 5 constrained inputs: 4 for blinds and 1 for heating
 - 7 modeled disturbances: 1 for occupancy, 2 for ambient and ground temperature, and 4 for solar radiation
- (iii) Large building, obtained from [37]
 - 113 states for temperatures, using comfort constraints between 20°C and 28°C when the building is occupied
 - 9 constrained inputs: 4 for blinds and 5 for heating
 - 11 modeled disturbances: 5 for occupancy, 2 for ambient and ground temperature, and 4 for solar radiation

We classify the buildings in residential (mainly occupied during nighttime) and commercial (mainly occupied during daytime). Usually, the buildings provide most reserves while being unoccupied, as they are less restricted then. For obtaining tractable bidding problems, we use linear decision rules as in Appendix A and [22, Section 4]. We use free licenses for Yalmip [41], Gurobi [42], Julia [43], and JuMP [44].

A. Example Scenario for Six Buildings

We sample six buildings from the test set and obtain an aggregated reserve bid with Alg. 2. Fig. 9 illustrates the resulting bids y_1 to y_6 (bottom), and the Lagrange multiplier Λ (top).

Fig. 9. Example scenario of six buildings placing a joint reserve offer by solving (6). The horizontal axis shows the bidding period. The day-night provision pattern identifies three residential (blue) and three commercial (red) buildings.

In Fig. 9, we show the bids $y_1, \dots, y_6$ as a fraction of the combined bid Y (not illustrated), which adds up to be time-constant. In a separate simulation, we solve the individual problem (4) for each building, and we compare the result to the solution of the aggregated problem as illustrated in Fig. 9. We observe that aggregated bidding with (6) yields a combined bid Y that is 57% larger than the accumulation of individual solutions from (4). This considerable aggregation advantage results from having a mixed group of residential and commercial buildings, which complement each other for satisfying the time-constant constraint. For the Lagrange multiplier, we see a day-night pattern that indicates a provision shortage at night. Hence, the commercial buildings are under-represented, and Lagrangian-based allocation places an incentive for more commercial buildings to join the group.

B. Analysis of the Necessary Number of ADMM Iterations

The feasible extraction scheme presented in Section III-D provides the opportunity to terminate the algorithm before it converges to a feasible solution. Therefore, the decision for algorithm termination only depends on the objective value

$$J(\{u_b(\cdot)\}_b, Y) := \sum_{b=1}^M c_b^\top u_b(0) - p^\top Y. \quad (16)$$

Fig. 10 shows how many ADMM iterations are necessary until the feasible-extraction objective $J^F := J(\{u_b^F(\cdot)\}_b, Y^F)$ comes within 1% of its precomputed optimal value $J^* > 0$.

Fig. 10. Required iterations of Alg. 2 for reaching $|J^F - J^*|/J^* \leq 1\%$ depending on the aggregation size M . For each aggregation size, we sample ten combinations of buildings from the building test set, and we show the average and range of the necessary number of ADMM iterations.

In Fig. 10, we observe: (i) on average only 18 ADMM iterations are needed to come within 1% of the best-achievable objective; (ii) the range between different instances of the same aggregation size is about ± 7 iterations; and (iii) the ADMM convergence characteristics (average and range) are fairly constant with the aggregation size. We expect the low number of iterations (i) and the small range (ii) as we have concentrated the main computational effort into the building-individual updates, which leaves a comparably simple and invariant task to the ADMM-based negotiation process. For the invariance to the aggregation size (iii), we note that the closed-form aggregation step (10) only depends on averaged information Ω . Hence, the mediation effort is not expected to vary with the number of buildings that contribute to this average. With feasible extraction and the observations from Fig. 10, we can be confident to use the simplest possible ADMM termination criterion, namely just stopping after a fixed number of iterations (e.g., 25). Besides being

conceptually simple, this stopping criterion also leads to a fixed number of communication rounds between the buildings, making the entire decision process well-timed and easy to handle. In particular, given that we impose a time-limit for solving the building-individual problem, the total negotiation time for obtaining a joint reserve bid is constant.

V. CONCLUSION

We discuss how flexible energy consumers, particularly buildings, can collaborate to provide frequency reserves in a reserve market setting. We show that when using a distributed ADMM-based approach, we obtain conceptual benefits, such as low set-up costs, fully-decentral computation, scalability, privacy, flexibility, and autonomy in the participating buildings. We also show that ADMM provides a reward allocation scheme that reduces provision imbalance and therefore places an incentive to strengthen the group. Finally, we emphasize the practicality of our approach by describing feasible extraction schemes, leading to a predictable stopping criterion and therefore a constant time for obtaining a joint reserve bid.

APPENDIX A

TRACTABLE FORMULATION OF BUILDING CONSTRAINTS

In [18], [22], three steps are taken to obtain a tractable reformulation of (5). First, the uncertain reserve request $s \in [-y, y]$ is represented through the normalized disturbance $\zeta = (\zeta^1; \dots; \zeta^N) \in [-1, 1] \subset \mathbb{R}^N$, such that $s = s(\zeta) = \text{diag}(y)\zeta$. This parametrization has the advantage that the normalized uncertainty set $[-1, 1]$ is independent of the decision variable y , which simplifies the maximization of the uncertainty set $[-y, y]$ [18, Section 3]. In a second step, the control policies $u(s)$ and $\Delta u(s)$ are reparametrized with respect to the normalized disturbance as $\tilde{u}(\zeta)$ and $\Delta \tilde{u}(\zeta)$, which is still non-restrictive [18, Prop. 3]. In a third (restrictive) step, affine parametrizations $\tilde{u}(\zeta) = K\zeta + \kappa$ and $\Delta \tilde{u}(\zeta) = F\zeta$ are used, where $K, F \in \mathcal{L} \subseteq \mathbb{R}^{N_m \times N}$, $\mathcal{L}$ is the set of lower block-diagonal matrices, and $\kappa \in \mathbb{R}^{N_m}$ is a vector. The optimization is then performed over K, F, κ and y , which results in

$$\begin{aligned} \mathcal{C} = \{ (u(\cdot), y) \mid & \exists (K, F, \kappa, \tilde{u}(\cdot), \Delta \tilde{u}(\cdot), x(\cdot)) \text{ s.t.} \\ & \tilde{u}(\zeta) = u(\text{diag}(y)\zeta) = K\zeta + \kappa, \quad \Delta \tilde{u}(\zeta) = F\zeta, \quad K, F \in \mathcal{L}, \\ & x(\zeta) = Ax^1 + B[\tilde{u}(\zeta) + \Delta \tilde{u}(\zeta)] + Ev, \\ & (x(\zeta), \tilde{u}(\zeta) + \Delta \tilde{u}(\zeta)) \in \mathcal{X} \times \mathcal{U}, \quad \forall \zeta \in [-1, 1], \\ & y = (I_N \otimes \eta^\top)F\zeta \} \end{aligned} \quad (17)$$

which is consistent with (5) in the case that $\mathcal{F}$ contains linear functions only. Note that (17) still has an infinite number of constraints. Following the procedure in [22, Section 4], it can be written as an exact dual reformulation with finitely many constraints. While the resulting problem is convex and finite-dimensional, its solution remains computationally challenging. The reason is that the parametrization of the input $\tilde{u}(\cdot)$ with the matrix K introduces a large number of decision variables. One way of reducing the computational complexity, at the cost of an increasingly conservative solution, is to impose structure on K . Popular strategies include restricting K to be block-diagonal, or even setting $K = 0$, which means that $\tilde{u}$ is pre-allocated, i.e., only $\Delta \tilde{u}$ depends on the actual reserve request.

APPENDIX B PROOFS

A. Proof of Proposition 1

Problem (7) is equivalent to (6). Finite-dimensionality of $\mathcal{C}$ allows for using the convergence result in [10]. We need the functions f, g to have a nonempty, closed, and convex epigraph, which is given by (9) and convexity of $\mathcal{C}_b$. Further, due to $\mathcal{C}_b \neq \emptyset$, there is a feasible solution with $y_b = Y = 0$. ■

B. Proof of Proposition 2

To simplify the notation, we use $y := (y_1; \dots; y_b)$ and similarly $\bar{y}$ and λ . The aggregation step then becomes

$$\min_{(Y^k, \bar{y})} \lambda^\top \bar{y} + \frac{\rho}{2} \|\bar{y} - y\|_2^2 - p^\top 1_{N \times 1} Y^k \quad (18a)$$

$$\text{s.t. } 1_{N \times 1} Y^k = (1_{1 \times M} \otimes I_N) \bar{y}, \quad (18b)$$

where $Y = 1_{N \times 1} Y^k$, i.e., Y is parametrized by a single decision variable $Y^k \in \mathbb{R}$, ensuring that $Y^1 = Y^2 = \dots = Y^N$. To avoid confusion with the solution of Alg. 2, we denote the solution to (18) as $(Y^k, \bar{y})$. Following [45, Section 16.2], this solution satisfies the KKT conditions $0 = 1_{1 \times N}(\eta + p)$, $\bar{y} = y + \frac{1}{\rho}(-\lambda - (1_{M \times 1} \otimes I_N)\eta)$, and $1_{N \times 1} Y^k = (1_{1 \times M} \otimes I_N) \bar{y}$, where η is the optimal value of the Lagrange multiplier associated to (18b). By using the Schur-complement method in [45], we show that the KKT system is solved by

$$Y^k = \frac{1}{\rho N} (1_{1 \times NM}(\rho y - \lambda) + M 1_{1 \times N} p) \quad (19a)$$

$$\eta = \frac{1}{M} ((1_{1 \times M} \otimes I_N)(\rho y - \lambda) - \rho 1_{N \times 1} Y^k) \quad (19b)$$

$$\bar{y} = \frac{1}{M} (M - 1_{M \times M} \otimes I_N)(y - \frac{1}{\rho} \lambda) + \frac{1}{M} I_{NM \times 1} Y^k, \quad (19c)$$

where we can write (19c) component-wise. Using the aggregate $\Omega = \frac{1}{M}(1_{1 \times M} \otimes I_N)(\rho y_b - \lambda_b)$ concludes the proof. ■

C. Proof of Proposition 3

Y^F satisfies the time-constant constraint (6d). From (12b) it is $\sum_{b=1}^M (y_b^k)^F = Y^F$, i.e., (6c) is satisfied. It remains to show that (12c) has a solution for all b . Due to $\sum_{b=1}^M y_b^k \geq (Y^k)^F$, equation (12b) ensures that $(y_b^k)^F \leq y_b^k$. Hence, by construction of $\mathcal{C}$ in (5), (12c) is always feasible. ■

D. Proof of Proposition 4

Inserting (10c) in step ③ of Alg. 2 yields $\lambda_b = \frac{\rho}{M} Y - \Omega$ for any b , which proves (ii). (i) follows as Y and Ω are aggregated quantities. In (iii), we consider solutions with $\bar{y}_b = y_b$, which we combine with (10b), (10c) to $\Lambda = \frac{1}{N} 1_{N \times N}(\Omega + p) - \Omega$. (iv) gives the same relation for Λ^F, Ω^F . In both cases, we multiply $1_{1 \times N}$ from the left and use $\frac{1}{N} 1_{1 \times N} 1_{N \times N} = 1_{1 \times N}$. The result is $1_{1 \times N} \Lambda^F = 1_{1 \times N} p$. (iii) and (iv) follow from writing $1_{1 \times N}$ as sum and multiplying Y^F from the right. ■

REFERENCES

- [1] Y. G. Rebours, D. S. Kirschen, M. Trotignon, and S. Rossignol, "A survey of frequency and voltage control ancillary services - part i: Technical features," *Power Systems, IEEE Transactions on*, vol. 22, no. 1, pp. 350–357, 2007.
- [2] Y. V. Makarov, C. Loutan, J. Ma, and P. de Mello, "Operational impacts of wind generation on California power systems," *Power Systems, IEEE Transactions on*, vol. 24, no. 2, pp. 1039–1050, 2009.
- [3] D. S. Callaway and I. A. Hiskens, "Achieving controllability of electric loads," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 184–199, Jan 2011.
- [4] F. Oldewurtel, T. Borsche, M. Bucher, P. Fortenbacher, M. G. Vaya, T. Haring, J. Mathieu, O. Mège, E. Vrettos, and G. Andersson, "A framework for and assessment of demand response and energy storage in power systems," in *Bulk Power System Dynamics and Control, Security and Control of the Emerging Power Grid*. IEEE, 2013, pp. 1–24.
- [5] M. A. Piette, J. Granderson, M. Wetter, and S. Kiliccote, "Responsive and intelligent building information and control for low-energy and optimized grid integration," in *ACEEE 2012 Summer Study on Energy Efficiency in Buildings*, Pacific Grove, CA, Oct 2012.
- [6] *Grundlagen Systemdienstleistungsprodukte*, Swissgrid AG, Feb 2017. [Online]. Available: https://www.swissgrid.ch/dam/swissgrid/experts/ancillary_services/Dokumente/D170214_AS-Products_V9R2_de.pdf
- [7] A. C. Chapman, G. Verbič, and D. J. Hill, "Algorithmic and strategic aspects to integrating demand-side aggregation and energy management methods," *IEEE Transactions on Smart Grid*, vol. 7, no. 6, pp. 2748–2760, 2016.
- [8] E. Vrettos, F. Oldewurtel, and G. Andersson, "Robust energy-constrained frequency reserves from aggregations of commercial buildings," *IEEE Transactions on Power Systems*, vol. 31, no. 6, pp. 4272–4285, 2016.
- [9] R. Glowinski and A. Marocco, "Sur l'approximation, par éléments finis d'ordre un, et la résolution, par pénalisation-dualité d'une classe de problèmes de dirichlet non linéaires," *Revue française d'automatique, informatique, recherche opérationnelle. Analyse numérique*, vol. 9, no. 2, pp. 41–76, 1975.
- [10] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.
- [11] M. Diekerhof, F. Peterssen, and A. Monti, "Hierarchical distributed robust optimization for demand response services," *IEEE Transactions on Smart Grid*, 2017.
- [12] E. Vrettos, F. Oldewurtel, F. Zhu, and G. Andersson, "Robust provision of frequency reserves by office building aggregations," *IFAC Proceedings Volumes*, vol. 47, no. 3, pp. 12 068–12 073, 2014.
- [13] H. Hao, A. Kowli, Y. Lin, P. Barooah, and S. Meyn, "Ancillary service for the grid via control of commercial building HVAC systems," in *American Control Conference*. Washington, USA: IEEE, 2013.
- [14] M. Maasoumy, B. Sanandaji, A. Sangiovanni-Vincentelli, and K. Poola, "Model predictive control of regulation services from commercial buildings to the smart grid," in *American Control Conference*. Portland, Oregon, USA: IEEE, 2014.
- [15] E. Vrettos and G. Andersson, "Scheduling and provision of secondary frequency reserves by aggregations of commercial buildings," *IEEE Transactions on Sustainable Energy*, vol. 7, no. 2, pp. 850–864, 2016.
- [16] C. Chen, J. Wang, Y. Heo, and S. Kishore, "MPC-based appliance scheduling for residential building energy management controller," *IEEE Transactions on Smart Grid*, vol. 4, no. 3, pp. 1401–1410, 2013.
- [17] E. Vrettos and G. Andersson, "Combined load frequency control and active distribution network management with thermostatically controlled loads," in *International Conference on Smart Grid Communications*. IEEE, Oct 2013, pp. 247–252.
- [18] X. Zhang, M. Kamgarpour, P. Goulart, and J. Lygeros, "Selling robustness margins: A framework for optimizing reserve capacities for linear systems," in *Conference on Decision and Control*. IEEE, 2014, pp. 6419–6424.
- [19] M. Balandat, F. Oldewurtel, M. Chen, and C. Tomlin, "Contract design for frequency regulation by aggregations of commercial buildings," in *Allerton Conference on Communication, Control, and Computing*, Sep 2014, pp. 38–45.
- [20] X. Zhang, E. Vrettos, M. Kamgarpour, G. Andersson, and J. Lygeros, "Stochastic frequency reserve provision by chance-constrained control of commercial buildings," in *European Control Conference*, Jul 2015, pp. 1134–1140.
- [21] T. T. Gorecki, A. Bitlisliolu, G. Stathopoulos, and C. N. Jones, "Guaranteeing input tracking for constrained systems: Theory and application to demand response," in *American Control Conference*. IEEE, Jul 2015, pp. 232–237.
- [22] X. Zhang, M. Kamgarpour, A. Georghiou, P. Goulart, and J. Lygeros, "Robust optimal control with adjustable uncertainty sets," *Automatica*, vol. 75, pp. 249–259, 2017.
- [23] A. Bitlislioglu, T. T. Gorecki, and C. N. Jones, "Robust tracking commitment," *IEEE Transactions on Automatic Control*, 2017.
- [24] E. Bilgin, M. C. Caramanis, I. C. Paschalidis, and C. G. Cassandras, "Provision of regulation service by smart buildings," *IEEE Trans. on Smart Grid*, vol. 7, no. 3, pp. 1683–1693, 2016.

- [25] A. F. Taha, N. Gatsis, B. Dong, A. Pipri, and Z. Li, "Buildings-to-grid integration framework," *IEEE Transactions on Smart Grid*, 2017.
- [26] M. Krating, E. Chu, J. Lavaei, S. Boyd *et al.*, "Dynamic network energy management via proximal message passing," *Foundations and Trends in Optimization*, vol. 1, no. 2, pp. 73–126, 2014.
- [27] S. Weckx and J. Driesen, "Load balancing with ev chargers and pv inverters in unbalanced distribution grids," *IEEE Transactions on Sustainable Energy*, vol. 6, no. 2, pp. 635–643, 2015.
- [28] X. Hou, Y. Xiao, J. Cai, J. Hu, and J. E. Braun, "Distributed model predictive control via proximal jacobian ADMM for building control applications," in *American Control Conference*. IEEE, 2017, pp. 37–43.
- [29] R. Verschae, H. Kawashima, T. Kato, and T. Matsuyama, "Coordinated energy management for inter-community imbalance minimization," *Renewable Energy*, vol. 87, pp. 922–935, 2016.
- [30] E. M. Burger and S. J. Moura, "Generation following with thermostatically controlled loads via alternating direction method of multipliers sharing algorithm," *Electric Power Systems Research*, vol. 146, pp. 141–160, 2017.
- [31] F. Oldewurtel, A. Parisio, C. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and M. Morari, "Use of model predictive control and weather forecasts for energy efficient building climate control," *Energy and Buildings*, vol. 45, pp. 15–27, 2012.
- [32] B. Lehmann, D. Gyalistras, M. Gwerder, K. Wirth, and S. Carl, "Intermediate complexity model for model predictive control of integrated room automation," *Energy and Buildings*, vol. 58, pp. 250–262, 2013.
- [33] F. Oldewurtel, "Stochastic model predictive control for energy efficient building climate control," Ph.D. dissertation, ETH Zurich, 2011.
- [34] X. Zhang, G. Schildbach, D. Sturzenegger, and M. Morari, "Scenario-based MPC for energy-efficient building climate control under weather and occupancy uncertainty," in *European Control Conference*, Jul 2013, pp. 1029–1034.
- [35] G. Darivianakis, A. Georghiou, R. Smith, and J. Lygeros, "A stochastic optimization approach to cooperative building energy management via an energy hub," in *Conference on Decision and Control*. Osaka, Japan: IEEE, Dec 2015, pp. 7814–7819.
- [36] P. J. Goulart, E. C. Kerrigan, and J. M. Maciejowski, "Optimization over state feedback policies for robust control with constraints," *Automatica*, vol. 42, no. 4, pp. 523–533, 2006.
- [37] G. Darivianakis, A. Georghiou, R. Smith, and J. Lygeros, "The power of diversity: Data-driven robust predictive control for energy efficient buildings and districts," *arXiv preprint arXiv:1607.05441*, 2016.
- [38] P. O. Scokaert and D. Mayne, "Min-max feedback model predictive control for constrained linear systems," *IEEE Transactions on Automatic Control*, vol. 43, no. 8, pp. 1136–1142, 1998.
- [39] J. Löfberg, *Minimax approaches to robust model predictive control*. Linköping University Electronic Press, 2003, vol. 812.
- [40] F. Rey, P. Hokayem, and J. Lygeros, "A tailored ADMM approach for power coordination in variable speed drives," in *Preprints of the 20th IFAC World Congress*, Jul 2017, pp. 7664–7669.
- [41] J. Löfberg, "Yalmip: A toolbox for modeling and optimization in matlab," in *Computer Aided Control Systems Design, 2004 IEEE International Symposium on*. IEEE, 2004, pp. 284–289.
- [42] I. Gurobi Optimization, "Gurobi optimizer reference manual," 2016. [Online]. Available: <http://www.gurobi.com>
- [43] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, "Julia: A fresh approach to numerical computing," *SIAM Review*, vol. 59, no. 1, pp. 65–98, 2017.
- [44] I. Dunning, J. Huchette, and M. Lubin, "JuMP: A modeling language for math. optimization," *SIAM Review*, vol. 59, no. 2, pp. 295–320, 2017.
- [45] J. Nocedal and S. J. Wright, *Numerical optimization*. Springer Science & Business Media, 1975, vol. 9, no. 4.

Felix Rey is a Ph.D. candidate in the Automatic Control Laboratory at ETH Zurich. He received his B.Eng. degree in Electrical Engineering and Information Technology from the University of Applied Sciences Constance (Germany, 2011), where he became a fellow of the German National Academic Foundation. He obtained his M.Sc. degree in Electrical Engineering and Information Technology from the Karlsruhe Institute of Technology (Germany, 2014). In 2014 he joined the Automatic Control Laboratory at ETH Zurich as doctoral student. His research interests include model predictive control, as well as distributed and embedded optimization, in particular with ADMM.

Xiaojing (George) Zhang received his B.Sc and M.Sc. degrees in electrical engineering and information technology, in 2010 and 2012, respectively, from the Swiss Federal Institute of Technology (ETH) Zurich, Switzerland. In 2016, he received his Ph.D. degree in automatic control, also from ETH Zurich. He is currently a post-doctoral fellow and the Associate Director at the Hyundai Center of Excellence at the University of California, Berkeley. His research interests include modeling, analysis, and control of stochastic uncertain systems, randomized and data-driven algorithms, as well as stochastic optimization, with application to autonomous vehicles, energy-efficient buildings and electric power systems.

Sandro Merkli received his BSc. and MSc. degrees in electrical engineering in 2012 and 2014, respectively, from ETH Zurich. He is currently a Ph.D. candidate at the Automatic Control Laboratory at ETH Zurich. His research interests are numerical optimization in general as well as its applications in power system operation and design.

Valentina Agliati received her BSc. and MSc. degrees in automation and control engineering in 2014 and 2017, respectively, from Politecnico di Milano. In 2016 she wrote her master's thesis in the Automatic Control Laboratory at ETH Zurich as an exchange student.

Maryam Kamgarpour is an Assistant Professor in the Automatic Control Lab at ETH Zurich. She obtained her Masters and Ph.D. in Control Systems at the University of California, Berkeley (2007, 2011) and her Bachelor of Applied Sciences in Systems Design Engineering from University of Waterloo, Canada (2005). Her research is on safety verification and optimal control of large-scale uncertain dynamical systems with applications in air traffic and power grid systems. She is the recipient of the NASA High Potential Individual Award, the NASA Excellence in Publication Award (2010) and the European Union (ERC) Starting Grant 2015.

John Lygeros received the B.Eng. degree in electrical engineering in 1990 and an M.Sc. degree in Systems and Control in 1991, both at Imperial College of Science Technology and Medicine, London, UK, in 1990 and 1991, respectively, and the Ph.D. degree from the Electrical Engineering and Computer Sciences Department, University of California, Berkeley, in 1996. He holds the chair of Computation and Control at the Swiss Federal Institute of Technology (ETH) Zurich, Switzerland, where he is currently serving as the Head of the Automatic Control Laboratory. During the period 1996–2000 he held a series of research appointments at the National Automated Highway Systems Consortium, Berkeley, the Laboratory for Computer Science, M.I.T., and the Electrical Engineering and Computer Sciences Department at U.C. Berkeley. Between 2000 and 2003 he was a University Lecturer at the Department of Engineering, University of Cambridge, UK, and a Fellow of Churchill College. Between 2003 and 2006 he was an Assistant Professor at the Department of Electrical and Computer Engineering, University of Patras, Greece. In July 2006 he joined the Automatic Control Laboratory at ETH Zurich, first as an Associate Professor, and since January 2010 as Full Professor. His research interests include modeling, analysis, and control of hierarchical, hybrid, and stochastic systems, with applications to biochemical networks, automated highway systems, air traffic management, power grids and camera networks. John Lygeros is a Fellow of the IEEE and a Member of the IET and of the Technical Chamber of Greece.