A two-layer decentralized approach to the optimal energy management of a building district with a shared thermal storage ⋆

Daniele Ioli * **Luca Deori** * **Alessandro Falsone** * **Maria Prandini** *

* *Dipartimento di Elettronica Informazione e Bioingegneria, Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133 Milano, Italy (e-mail: {daniele.ioli, luca.deori, alessandro.falsone, maria.prandini}@polimi.it)*

Abstract: We propose a decentralized scheme for the energy management of a district composed of multiple buildings. Each building is equipped with its own chiller, whereas a single thermal storage is shared among all buildings. The buildings aim at jointly minimizing the electric energy costs while guaranteeing comfort conditions for their occupants, and they can act on their own temperature set-point and the usage of the common storage to this purpose. The problem can be formulated as a constrained optimization program where the buildings decision variables are coupled via some global constraints, due to the shared storage, and the cost function, since the electric energy price depends on the district demand. To distribute the computational load and reduce the amount of transmitted information, we propose a two-layer solution where a central entity takes care of the global constraints by updating some dual variables (outer layer), whereas the primal variables, i.e., zone temperature set-point and storage usage, are optimized locally by the buildings through nested iterations where the price is recomputed based on the updated information on the district demand provided by the central entity (inner layer).

Keywords: Decentralized optimization, cooperative multi-agent system, energy management operations, smart grid, thermal storage sharing.

1. INTRODUCTION

We address optimal energy management for the thermal control of a district network with multiple buildings sharing a storage unit. Energy storage systems add flexibility to energy management systems (Wang et al. (2013)), and are largely studied as a means for compensating the intermittent nature of renewable energy sources (see e.g., Teleke et al. (2010); Telaretti et al. (2015)). However, they are typically expensive resources and rarely employed at their full capacity when they are private resources of single users. This motivates our setting, where a single storage is shared among multiple buildings.

In our set-up, each building is equipped with a chiller plant that converts electrical energy into cooling energy and is operated so as to maintain an adequate room temperature and guarantee comfort conditions. The thermal storage can act as an energy buffer, allowing load request deferral, and giving the possibility to operate the chiller plants at their maximum efficiency and to buy electrical energy from the grid when the price is lower. These advantages are further enhanced by the fact that the thermal storage unit is shared: buildings with a chiller plant that has a limited capability – and may not be able to satisfy load picks – could draw some additional cooling energy from the storage, recharged by some larger chiller plant for which it may be even convenient to work at a higher production rate in order to improve its efficiency. As a result, buildings would end up sharing their actuation systems (their chillers), thus realizing in practice a spatially distributed thermal control system that is more robust to local picks in load request, Belluschi et al. (2016).

Appropriate energy management strategies need to be conceived so as to establish how the buildings can share the storage unit. Here, we suppose that buildings are cooperating aiming at minimizing the electric energy cost for the thermal control of the district, and that the energy price is a time-varying affine function of the district demand: the higher the demand, the higher is the price. Similar price structures are adopted e.g. in Ma et al. (2013); Deori et al. (2016,7); Gharesifard et al. (2016); Grammatico et al. (2016) addressing charging control of plug-in electric vehicles. In our context, we can think of the grid operator applying a pricing strategy to level the electric energy request. The interested reader is referred to Gharesifard et al. (2016) for some discussion on pricing strategies within the energy market. Some further flexibility is added to the energy management system by allowing each building to slightly modulate the temperature set-points of its thermally controlled zones within some comfort bounds. The energy management problem can then be formulated as a constrained optimization program where the sum of the electric energy costs of all buildings is minimized with respect to the zone temperature set-points and the exchanges with the storage as optimization variables, subject to local and global constraints.

The size of the optimization problem grows linearly with the number of buildings, and its centralized solution would

⋆ This research was supported by the European Commission under the project UnCoVerCPS, grant number 643921.

require to collect various pieces of information that are intrinsically local, like cooling energy requests, zone temperature set-points, chiller characteristics, comfort conditions that depend on the building usage (e.g., commercial or residential). To overcome these computational and communication issues, we propose a decentralized iterative solution where only part of these pieces of information need to be conveyed to some central entity. A numerical example shows the performance of the proposed decentralized two-layer scheme versus an alternative heuristic approach where buildings are assigned a fixed share of the storage.

Building energy management, and climate control in particular, has recently attracted the attention of various researchers (see, e.g., Siroky et al. (2011); Ma et al. (2012); Deng et al. (2013); Vignali et al. (2017); Scherer et al. (2014)). Indeed, energy consumption in buildings represents approximately 40% of the worldwide energy demand, and more than half of this amount is spent for climate control (Pérez-Lombard et al. (2008)). Optimization based methods are largely adopted to address energy management in buildings, since the problem is naturally formulated as the optimization of some objective function (e.g., energy costs, deviation from some desired consumption) subject to physical and technological constraints (e.g., actuation limits, comfort constraints). In all the aforementioned references the underlying algorithm refers to a single building or room. In Ioli et al. (2015a), a hierarchical scheme implementing a decentralized heuristic solution is proposed. In Chang et al. (2013) decentralized control is applied to a home energy management problem, whereas in Kranning et al. (2014) the Alternating Direction Method of Multipliers (ADMM, Boyd et al. (2010)) is adopted for decentralized optimal power flow in transmission networks. A distributed approach is proposed in Belluschi et al. (2016), within our same modeling framework. The proposed solution does not require any central entity, is resilient to (temporary) failures in the communications, and preserves the private information on building electricity consumption. However, it applies only to the case when the cost function is separable, which is not the one addressed in this work where the electric energy price is dependent on the district demand.

2. PROBLEM FORMULATION

We consider a district network composed of m buildings. Each building is equipped with a chiller plant and is divided into some thermally conditioned zones, each one characterized by its own (average) temperature. Let T_z^i be the vector of set-points for zones temperatures relative to building i. Our objective is to determine an optimal profile for T_z^i, $i = 1, \dots, m$, along some finite time-horizon $[t_0, t_f]$, so as to minimize the electric energy cost for cooling the district, while guaranteeing a satisfactory level of comfort to the occupants of each building. To this end, we adopt the modeling approach in Ioli et al. (2015b), where the time horizon $[t_0, t_f]$ is discretized into M time slots of length dt and the profile for T_z^i, $i = 1, \dots, m$, is assumed to be linear in each time slot with a slope limited by the capability of the chiller plant so that it can actually be tracked. Let us denote as $T_z^i(k)$ the temperature set-points at the end of the k-th time slot and $T_z^i(0)$ the zone temperatures at t_0 of building i.

Following Ioli et al. (2015b), the electrical energy consumption $E_\ell^i(k)$ of building i is expressed as a differentiable convex function of the energy $E_{ch}^i(k)$ requested to the chiller of the i-th building during the time slot k, namely

$$E_\ell^i(k) = c_1^i E_{ch}^i(k)^4 + c_2^i E_{ch}^i(k)^2 + c_3^i, \tag{1}$$

where c_1^i, c_2^i, and c_3^i are appropriate coefficients characterizing the chiller of the i-th building.

Buildings in the district are sharing a storage unit, which allows to shift in time the electric energy requests so as to operate the chillers at higher efficiency or buy energy when the price is lower. Let $S(k)$ denote the amount of energy accumulated in the thermal storage at the end of the k-th time slot, $S(0)$ being its initial energy content. Then, the storage dynamics can be described by a first-order dynamical system through the recursive equation

$$S(k+1) = aS(k) - \sum_{i=1}^{m} s^i(k), \tag{2}$$

where $s^i(k)$ is the amount of energy drawn from ($s^i(k) > 0$) or released to ($s^i(k) < 0$) the storage by building i during time slot k, and $a \in (0,1)$ is a coefficient which models energy losses. The amount of energy requested to the chiller of the i-th building during the k-th time slot can then be computed as

$$E_{ch}^i(k) = E_c^i(k) - s^i(k), \tag{3}$$

where $E_c^i(k)$ is the amount of cooling energy requested by building i during time slot k to track the zone temperature profile in that time slot.

The energy cost along $[t_0, t_f]$ is computed as

$$\mathcal{C} = \sum_{k=1}^{M} \psi(k) \sum_{i=1}^{m} E_\ell^i(k), \tag{4}$$

where $\psi(k)$ is the unitary energy cost within the k-th time slot. In each time slot k, the energy price is taken to be the sum of two contribution: i) a baseline term $\beta(k)$ and ii) an additional term which is proportional to the energy demand in the district, i.e.,

$$\psi(k) = \beta(k) + \alpha(k) \sum_{i=1}^{m} E_\ell^i(k), \tag{5}$$

$\alpha(k)$ being a suitable (positive) coefficient. Through the choice of $\beta(k)$, the adopted price model is able to capture jointly the minimum price of the electrical energy and the effect on the price of other electric energy requests.

Let us define the following vectors $\boldsymbol{T_z^i} = [T_z^i(0) \cdots T_z^i(M)]^\top$, $\boldsymbol{s^i} = [s^i(1) \cdots s^i(M)]^\top$, $\boldsymbol{E_c^i} = [E_c^i(1) \cdots E_c^i(M)]^\top$, $\boldsymbol{E_{ch}^i} = [E_{ch}^i(1) \cdots E_{ch}^i(M)]^\top$, $\boldsymbol{E_\ell^i} = [E_\ell^i(1) \cdots E_\ell^i(M)]^\top$, $\boldsymbol{S} = [S(0) \cdots S(M)]^\top$ for all buildings $i = 1, \dots, m$. The storage energy content can then be expressed as

$$\boldsymbol{S} = \Xi_0 S(0) - \Xi_1 \sum_{i=1}^{m} \boldsymbol{s^i}, \tag{6}$$

where Ξ_0 and Ξ_1 are appropriate matrices. Furthermore, according to the building model in Ioli et al. (2015b), $\boldsymbol{E_c^i}$ can be expressed as $\boldsymbol{E_c^i} = \Omega_0^i \boldsymbol{T_z^i} + \Omega_1^i$, where Ω_0^i and Ω_1^i are appropriate matrices that depend on the i-th building structure, initial state, and external disturbance profiles, which are here assumed as given. The cooling energy request to chiller i is then given by $\boldsymbol{E_{ch}^i} = \Omega_0^i \boldsymbol{T_z^i} + \Omega_1^i - \boldsymbol{s^i}$

and $\boldsymbol{E_\ell^i}$ obtained by plugging $\boldsymbol{E_{ch}^i}$ into the differentiable convex function (1) is convex and differentiable with respect to the optimization variables $\boldsymbol{T_z^i}$ and $\boldsymbol{s^i}$.

The optimal energy management problem for the district network can then be formalized as

$$\min_{\{\boldsymbol{T_z^i},\boldsymbol{s^i}\}_{i=1}^m} \quad J(\boldsymbol{T_z^1},\dots,\boldsymbol{T_z^m},\boldsymbol{s^1},\dots,\boldsymbol{s^m}) \tag{7}$$

subject to:

$$\begin{cases} S_{\min} \le \boldsymbol{S}(\boldsymbol{s^1},\dots,\boldsymbol{s^m}) \le S_{\max} \\ e_M^\top \boldsymbol{S}(\boldsymbol{s^1},\dots,\boldsymbol{s^m}) \ge S(0) \end{cases} \tag{8}$$

$$\begin{cases} |\boldsymbol{s^i}| \le s_{\max}^i \\ \boldsymbol{E_c^i}(\boldsymbol{T_z^i}) \ge 0 \\ 0 \le \boldsymbol{E_{ch}^i}(\boldsymbol{T_z^i},\boldsymbol{s^i}) \le E_{ch,\max}^i \\ \boldsymbol{T_{z,\min}^i} \le \boldsymbol{T_z^i} \le \boldsymbol{T_{z,\max}^i} \\ T_z^i(M) = T_z^i(0) = T_{z,t_0}^i \end{cases} \quad i=1,\dots,m \tag{9}$$

where (8) are the coupling constraints and (9) are the constraints that apply locally to each building. For the sake of clarity, the dependence of all quantities from the optimization variables are made explicit in (7). In the problem formulation, $S_{\min}$ and $S_{\max}$ are the minimum and maximum amount of energy that can be stored and, for each building i, $\boldsymbol{T_{z,\min}^i}$ and $\boldsymbol{T_{z,\max}^i}$ represent the comfort constraints, $E_{ch,\max}^i$ is the maximum amount of cooling energy that chiller i can provide during each time slot, $s_{\max}^i$ is the maximum amount of energy that can be exchanged with the common storage per time slot. The constraints $\boldsymbol{E_{ch}^i} \ge 0$ and $\boldsymbol{E_c^i} \ge 0$ ensure that no heating energy is requested, while $S(M) = e_M^\top \boldsymbol{S} \ge S(0)$ avoid the storage depletion at the end of the time horizon, e_M being the M-dimensional column vector with all zeros except for a one in the M-th position. The constraint $T_z^i(0) = T_{z,t_0}^i$ sets the initial temperature of the zones to their actual temperature at t_0, whilst $T_z^i(M) = T_z^i(0)$ enforces the final temperature to be identical to the initial one in view of a repetitive application of the energy management strategy e.g. on a one-day time basis. Finally, the objective function in (7) is given by

$$J = \mathcal{C} + \sum_{i=1}^m \left(\rho_s \|\boldsymbol{s^i}\|^2 + \rho_T \|\boldsymbol{T_z^i} - \overline{\boldsymbol{T}}\boldsymbol{_z^i}\|^2\right), \tag{10}$$

which is the sum of the electric energy cost $\mathcal{C}$ in (4) and two terms that penalize the control effort by limiting the extent of the energy exchanges with the storage and the deviation of the zone temperature set-points from some nominal profile $\overline{\boldsymbol{T}}\boldsymbol{_z^i}$. The coefficients $\rho_s > 0$ and $\rho_T > 0$ set the relative importance between the electric energy cost and the control effort.

We shall next show strict convexity and differentiability of $J(\cdot)$, which are necessary conditions for the application and technical soundness of the methodology proposed in the next section. To this purpose, note first that, by substituting (5) in (4), $\mathcal{C}(\cdot)$ can be expressed as

$$\mathcal{C} = \sum_{k=1}^M \left[\beta(k)\left(\sum_{i=1}^m E_\ell^i(k)\right) + \alpha(k)\left(\sum_{i=1}^m E_\ell^i(k)\right)^2\right]. \tag{11}$$

For any $k = 1,\dots,M$, $E_\ell^i(k)$ is a convex function of the optimization variables $\boldsymbol{T_z^i}$ and $\boldsymbol{s^i}$, $i=1,\dots,m$. Then, also $\sum_{i=1}^m E_\ell^i(k)$ is convex. Moreover, the fact that $\sum_{i=1}^m E_\ell^i(k)$ is always non-negative together with the monotonicity property of the square function for non-negative arguments, leads to the convexity of $\left(\sum_{i=1}^m E_\ell^i(k)\right)^2$. Finally, $\alpha(k)$ and $\beta(k)$ are positive coefficients for all $k = 1,\dots,M$, therefore $\mathcal{C}(\cdot)$ is convex in $\boldsymbol{T_z^1},\dots,\boldsymbol{T_z^m},\boldsymbol{s^1},\dots,\boldsymbol{s^m}$. Due to convexity of $\mathcal{C}$ and strict convexity of the penalization terms in (10), $J(\cdot)$ is strictly convex. Differentiability of $\mathcal{C}(\cdot)$ as a function of the optimization variables $\boldsymbol{T_z^i}$ and $\boldsymbol{s^i}$, $i=1,\dots,m$ follows from the fact that it is differentiable as a function of its argument $E_\ell^i(k)$, which, in turn, is differentiable as a function of $\boldsymbol{T_z^i}$ and $\boldsymbol{s^i}$, $i=1,\dots,m$, for any $k=1,\dots,M$. Differentiability of $J(\cdot)$ then follows from its definition in (10).

3. PROPOSED DECENTRALIZED SOLUTION

By analyzing (7), one can see that each building has its own optimization variables and set of local constraints, but its decision is coupled to that of the others through the constraints (8) on the energy accumulated in the shared storage and through the cost function (10), since each building demand affects the electric energy price. As the number of buildings inside the district increases, the computational effort involved in a centralized solution to (7) grows rapidly. Additionally, various pieces of information that are intrinsically local, and possibly private, need to be collected. We here seek for a decentralized approach that addresses the computation and communication issues of a centralized solution while preserving optimality. The proposed approach integrates the methods described next that address the coupling due to the constraints (8) and the cost function (10), respectively.

3.1 Handling the coupling constraints

This method rests on strong duality theory and allows to account for the coupling constraints (8) in (7), separately from the other local constraints (9), while preserving feasibility (and optimality) of the solution thanks to the problem convexity. By replacing $\boldsymbol{S}$ with its expression as a function of $\boldsymbol{s^i}$ in (6), the three coupling constraints in (8) can be rewritten in compact form as $A\sum_{i=1}^m \boldsymbol{s^i} - b \le 0$, for appropriately defined matrices A and b. According to duality theory, the dual of (7) obtained by dualizing the coupling constraints is given by

$$\max_{\lambda \ge 0} \min_{\{\boldsymbol{T_z^i},\boldsymbol{s^i}\}_{i=1}^m} \quad J + \lambda^\top\left(A\sum_{i=1}^m \boldsymbol{s^i} - b\right) \tag{12}$$

subject to:

$$\begin{cases} |\boldsymbol{s^i}| \le s_{\max}^i \\ \boldsymbol{E_c^i} \ge 0 \\ 0 \le \boldsymbol{E_{ch}^i} \le E_{ch,\max}^i \\ \boldsymbol{T_{z,\min}^i} \le \boldsymbol{T_z^i} \le \boldsymbol{T_{z,\max}^i} \\ T_z^i(M) = T_z^i(0) = T_{z,t_0}^i \end{cases} \quad i=1,\dots,m, \tag{13}$$

where λ is the vector of Lagrange multipliers associated with the coupling constraints only.

If we define the optimization variables $x^i = [\boldsymbol{T_z^i}^\top \ \boldsymbol{s^i}^\top]^\top$ for building i and introduce its local constraint set X^i as originated by all (local) constraints (13), then problem (12) can be rewritten as

$$\max_{\lambda \ge 0} \min_{\{x^i \in X^i\}_{i=1}^m} f(x^1,\dots,x^m,\lambda), \tag{14}$$

where the function

$$f(x^1,\ldots,x^m,\lambda) = J + \lambda^\top\big(A\sum_{i=1}^m \boldsymbol{s^i} - b\big) \tag{15}$$

depends on all optimization variables of all buildings and not only on the storage exchanges $\boldsymbol{s^i}$ because $J = J(\boldsymbol{T_z^1},\ldots,\boldsymbol{T_z^m},\boldsymbol{s^1},\ldots,\boldsymbol{s^m})$. The dual ascent algorithm (Bertsekas (1999)) can provide a numerical solution to (14) by iteratively performing an update of the primal variables $x = [(x^1)^\top \cdots (x^m)^\top]^\top$, and an update of the dual variables λ. Each iteration κ_d of the dual ascent algorithm consists of two update steps: in the first one, called *primal update*, the primal variables x are updated solving an optimization problem where the cost $f(\cdot)$ is minimized with respect to x, accounting for constraints X^i, $i = 1,\ldots,m$, while dual variables are kept fixed to the value at the current iteration $\lambda = \lambda(\kappa_d)$:

$$x(\kappa_d+1) \in \arg\min_{\{x^i\in X^i\}_{i=1}^m} f(x^1,\ldots,x^m,\lambda(\kappa_d)). \tag{16}$$

In the second step, called *dual update*, the dual variables λ are updated performing a projected gradient step

$$\lambda(\kappa_d+1) = \big[\lambda(\kappa_d) + \delta(\kappa_d)\big(A\sum_{i=1}^m \boldsymbol{s^i}(\kappa_d) - b\big)\big]^+, \tag{17}$$

where $\delta(\kappa_d) > 0$ is the step-size and $[\,\cdot\,]^+$ denotes the projection operator onto the non-negative orthant to which λ is constrained ($\lambda \geq 0$). Due to strict convexity of the primal objective function, the dual function is continuously differentiable everywhere (Bertsekas, 1999, Proposition 6.1.1) but it is not guaranteed to have a Lipschitz continuous gradient. Therefore, in order for the sequence $\{\lambda(\kappa_d)\}_{\kappa_d\geq 0}$ to converge to an optimal dual solution, the step-size $\delta(\kappa_d)$ cannot be kept constant but it needs to asymptotically vanish at an appropriate rate, i.e., $\lim_{\kappa_d\to\infty}\delta(\kappa_d) = 0$ and $\sum_{\kappa_d=0}^{\infty}\delta(\kappa_d) = +\infty$. Moreover, due to strict convexity of the primal objective function $J(\cdot)$, the sequence $\{x(\kappa_d)\}_{\kappa_d\geq 0}$ converges to an optimal solution of (7) (Boyd et al. (2010)).

3.2 Handling the coupling cost function

Note that the primal update step (16) (with $\lambda(\kappa_d)$ fixed) has an almost separable structure in that each building has its own local constraint set X^i and the coupling is only due to the objective function $f(\cdot)$. We can then adopt for its solution the decentralized iterative algorithm developed in Deori et al. (20167) and briefly recalled next.

Consider the optimization problem

$$\min_{\{\zeta^i\in Z^i\}_{i=1}^m} \varphi(\zeta^1,\ldots,\zeta^m). \tag{18}$$

Each agent i starts with a tentative value $\zeta_0^i \in Z^i$ for the optimal solution to (18), and send this information to the central entity. At the generic (κ_p+1)-th iteration, each agent i receives from the central entity a vector ζ_{κ_p} which contains the solution of all agents at the previous iteration, and updates its current estimate of the local variables ζ^i performing the minimization step

$$\zeta^i_{\kappa_p+1} = \arg\min_{\zeta^i\in Z^i}\big\{\varphi(\zeta^i,\zeta^{-i}_{\kappa_p}) + c\|\zeta^i - \zeta^i_{\kappa_p}\|^2\big\}, \tag{19}$$

where ζ^{-i} denotes the vector emanating from $\zeta = (\zeta^1,\ldots,\zeta^m)$ when ζ^i is removed. Finally, it sends the updated estimate to the central authority. The idea of (19) is that each agent i optimizes with respect to its own decision variables $\zeta^i \in Z^i$, the weighted sum of two contributions: the cost function $\varphi(\cdot,\zeta^{-i}_{\kappa_p})$, where the variables of all other agents are fixed to the value at the previous iteration $\zeta^{-i}_{\kappa_p}$, and a proximal term that penalizes the distance of the new estimate from the previous one $\zeta^i_{\kappa_p}$. The relative importance between the two terms is dictated by the weight $c > 0$. If $c > \frac{m-1}{2m-1}\sqrt{m}L$, L being the Lipschitz constant of the gradient $\nabla\varphi(\cdot)$ of $\varphi(\cdot)$ as a function of ζ, then, convergence is guaranteed (Deori et al. (20167)).

We can then apply (19) iteratively to perform the primal update (16) of the dual ascent algorithm in a decentralized fashion by setting $\zeta^i = x^i$, $Z^i = X^i$, and $\varphi(\cdot) = f(\cdot,\lambda(\kappa_d))$ at every iteration κ_d. More precisely, let $\tau^i_{\kappa_p}$ and $\sigma^i_{\kappa_p}$ be the $\boldsymbol{T_z^i}$ and $\boldsymbol{s^i}$ components of the solution $\xi_{\kappa_p} = [\tau^{i\top}_{\kappa_p}\ \sigma^{i\top}_{\kappa_p}]^\top$ to (19) at the inner iteration $\kappa_p - 1$. Define as $\boldsymbol{E^i_{\ell}}_{\kappa_p}$ vector $\boldsymbol{E^i_\ell}$ evaluated at $\boldsymbol{T^i_z} = \tau^i_{\kappa_p}$ and $\boldsymbol{s^i} = \sigma^i_{\kappa_p}$. Writing (19) at iteration κ_p with $f(x^i,\xi^{-i}_{\kappa_p},\lambda(\kappa_d))$ in place of $\varphi(\cdot)$, x^i in place of ζ^i, and $\xi^i_{\kappa_p}$ in place of $\zeta^i_{\kappa_p}$, after some simple computations, we get

$$\begin{aligned}\xi^i_{\kappa_p+1} = \arg\min_{x^i\in X^i}\Big\{&\mathcal{C}\big(x^i,\textstyle\sum_{j\neq i}\boldsymbol{E^j_{\ell}}_{\kappa_p}\big) + \lambda^\top(\kappa_d)A\boldsymbol{s^i} \\ &+ \rho_s\|\boldsymbol{s^i}\|^2 + \rho_T\|\boldsymbol{T^i_z} - \overline{\boldsymbol{T}}^{\boldsymbol{i}}_{\boldsymbol{z}}\|^2 + c\|x^i - \xi^i_{\kappa_p}\|^2\Big\}.\end{aligned} \tag{20}$$

3.3 Two-layer algorithm

We can now integrate the methods described in Sections 3.1 and 3.2 in a two-layer nested approach, where step (17) of the dual ascent constitutes the outer layer and is performed by a central entity, whereas step (16) is implemented through an inner layer involving a loop where the central entity and the buildings are jointly performing the decentralized iterative approach in Section 3.2. This results in Algorithm 1, where only steps 9-11 and 16 are meant to be performed (in parallel) by the buildings. Buildings need to exchange with the central entity their electric energy consumption profiles and the shared storage usage. This avoids overloading the communication links by transmitting additional (private) pieces of information. The central entity verifies if all solutions are within their thresholds and decides whether to stop the inner/outer loop or not. Note that as for the inner loop, convergence can be assessed by the central entity without requiring the buildings to transmit their local solutions. It is indeed sufficient that each building checks locally if changes in its solution $\xi^i_{\kappa_p}$ are within a certain threshold and sends a binary information to the central entity.

Note that (7) is a sharing problem that can be solved using ADMM (see (Boyd et al., 2010, Section 7.3)). However, the central entity would need to solve an optimization problem, whereas in Algorithm 1 it is only required to update and store the dual variables (step 17) and an aggregate information on the overall consumption (step 13). Furthermore, when the coupling in the cost (11) (square term) and the coupling in the constraints (8) ($\boldsymbol{S}$ term) do not depend on the sum of optimization variables, then (7) is not a sharing problem anymore and the general form of ADMM ((Boyd et al., 2010, Section 3.1)) has to be applied. This entails that the central entity would need to solve an optimization problem that scales linearly with the number of agents and is almost as difficult as the original one.

Algorithm 1 Two-layer algorithm

% Outer layer initialization
$\kappa_d \leftarrow 0$; central entity chooses $\lambda(0) \geq 0$
repeat *% Outer layer*
  % Inner layer initialization
  $\kappa_p \leftarrow 0$; buildings choose $\xi_0 \in X^i$, $i = 1, \ldots, m$
  $\boldsymbol{E}_{\boldsymbol{\ell}0}^{tot} \leftarrow \sum_{i=1}^{m} \boldsymbol{E}_{\boldsymbol{\ell}0}^{\boldsymbol{i}}$
  repeat *% Inner layer*
    for $i = 1$ **to** m **all buildings in parallel do**
      receive $\boldsymbol{E}_{\boldsymbol{\ell}_{\kappa_p}}^{tot}$ from the central entity
      compute $\xi^i_{\kappa_p+1} \leftarrow$ (20)
      send $\boldsymbol{E}^{\boldsymbol{i}}_{\boldsymbol{\ell}_{\kappa_p+1}}$ to central entity
    end for
    $\boldsymbol{E}_{\boldsymbol{\ell}_{\kappa_p+1}}^{tot} \leftarrow \sum_{i=1}^{m} \boldsymbol{E}^{\boldsymbol{i}}_{\boldsymbol{\ell}_{\kappa_p+1}}$ central entity update
    $\kappa_p \leftarrow \kappa_p + 1$
  until a stopping criterion on ξ_{κ_p} is met
  Buildings send $\sigma^i_{\kappa_p}$, $i = 1, \ldots, m$, to central entity
  $\lambda(\kappa_d + 1) = \left[\lambda(\kappa_d) + \delta(\kappa_d)\left(A \sum_{i=1}^{m} \sigma^i_{\kappa_p} - b\right)\right]^+$
  $\kappa_d \leftarrow \kappa_d + 1$
until a stopping criterion on $\lambda(\kappa_d)$ is met

4. NUMERICAL EXAMPLE

We considered a district with $m = 3$ commercial buildings that jointly optimize the electrical energy cost over a one-day horizon discretized in $M = 144$ time slots of $dt = 10$ minutes each. The buildings have the same structure, which can be found in Ioli et al. (2015b), and are equipped with different chiller plants described by (1) with the following sets of coefficients $(c_1^1, c_2^1, c_3^1) = (3.42 \cdot 10^{-4}, 3.69 \cdot 10^{-2}, 1.46)$, $(c_1^2, c_2^2, c_3^2) = (5.21 \cdot 10^{-5}, 2.16 \cdot 10^{-2}, 2.82)$, and $(c_1^3, c_2^3, c_3^3) = (5.17 \cdot 10^{-6}, 1.49 \cdot 10^{-2}, 5.22)$ where c_1^i, c_2^i, and c_3^i are measured in MJ^{-3}, MJ^{-1}, and MJ, respectively, and with maximum cooling energy production per time slot given by $E^1_{ch,\max} = 16$, $E^2_{ch,\max} = 30$, and $E^3_{ch,\max} = 40$ in MJ. In the sequel we shall then refer to chiller $i = 1, 2, 3$ as "small", "medium", and "large" chiller, respectively. We consider a single zone per building. During working hours (from 8AM to 6PM) the zone temperature $T_z^i(k)$ of each building is required to be in the range 22-24°C. The range is extended to 16-30°C for other time slots. The initial temperature is set equal to $T^i_{z,t_0} = 24$°C for all buildings. The shared thermal storage minimum and maximum capacity are $S_{\min} = 75$MJ and $S_{\max} = 1425$MJ (5% and 95% of the overall capacity 1500MJ). The rate of thermal energy exchange per time slot is upper bounded by $s^i_{\max} = 11$MJ, $i = 1, 2, 3$. The initial storage content is $S(0) = 750$MJ. The storage coefficient $a = 0.9983$ in (2) models a 1% energy loss per hour. In the price expression (5), $\beta(k)$ is the time-varying energy price in Ioli et al. (2015b), whereas $\alpha(k)$ is constant and equal to 0.0023, which is 10% of the average of $\beta(k)$ over $[1, M]$. As for the cost function J in (10), $\overline{\boldsymbol{T}}^{\boldsymbol{i}}_{\boldsymbol{z}}$ is set constant and equal to 23°C for all buildings, and the coefficients $\rho_T = 10^{-2}$ and $\rho_s = 10^{-3}$ make both penalization terms of the same order of magnitude of the energy cost $\mathcal{C}$. The step size of the outer layer is set to $\delta(\kappa_d) = \bar{\delta}/(\kappa_d + 1)^\gamma$ with $\bar{\delta} = 10^{-6}$ and $\gamma = 0.1$. As for c in the inner layer, according to Deori et al. (20167), one would need to choose c so as to satisfy $c > \frac{m-1}{2m-1}\sqrt{m}L$, where L is the Lipschitz constant of $\nabla\varphi(\cdot) = \nabla f(\cdot, \lambda(\kappa_d))$ in (15). Computing L exactly is however difficult and (conservative) estimates of L lead to a very high value of c that slows down convergence excessively. We then set $c = 1$, which was experimentally verified to be appropriate in our numerical example. We chose as stopping criterion for both layers the condition that the norm of the relative or absolute difference between two subsequent iterations of the monitored quantities gets lower than 10^{-3}.

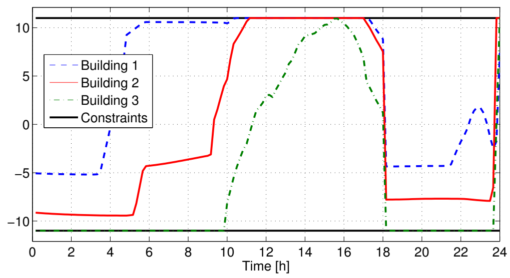

Fig. 1. Storage energy exchange [MJ]: optimal solution.

To asses the performance of the proposed methodology we compare it against the heuristic solution obtained when the storage is equally partitioned among the buildings, and each building has access only to its share, with an initial energy content equal to half of it. In this case the problem does not have any coupling constraints, and one can directly apply the decentralized algorithm in Section 3.2. Resulting costs are reported in Table 1, together with those of the solution without storage, thus pointing out the main role of the storage in cost saving. In Figure 1 we report the energy exchanges between the buildings and the thermal storage, which reveal that the building with the large chiller is providing thermal energy to the building with the small chiller through the shared thermal storage. The presence of the joint storage can indeed compensate for an imbalance between chiller size and building demand, thus allowing a sharing of the actuation capabilities in the district. Figure 2 reports the Coefficient Of Performance (COP), i.e., the ratio between cooling energy production and electric energy consumption per time slot, of the three chillers for both the optimal and heuristic solutions. By inspection, we can clearly see how the flexibility introduced by the shared thermal storage is exploited to make the small and large chillers operate at higher efficiency. The optimal temperature set-points are plotted in Figure 3, which shows that building 1 is subject to a stronger pre-cooling phase than the other buildings. To get an idea on the extent of the disturbances and their impact on the thermal behavior of the buildings, in Figure 3 we also plot the evolution of their temperatures without cooling.

Table 1. Energy cost $\mathcal{C}$ and cost function J.

	Optimal	Heuristic	No storage
$\mathcal{C}$	148.79	203.00	427.19
J	156.86	210.35	429.82

5. CONCLUSIONS AND FUTURE WORK

This paper deals with decentralized energy management for the thermal regulation of a district network composed of multiple building sharing a thermal storage. A main

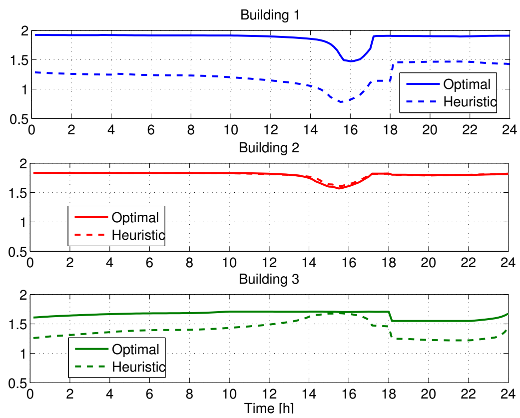

Fig. 2. COP profiles: optimal and heuristic solutions.

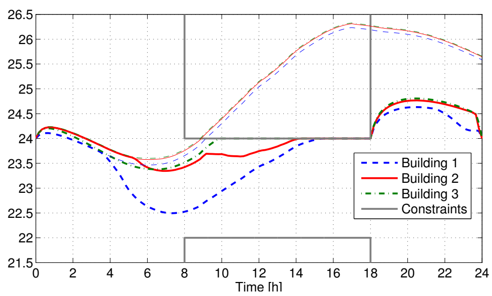

Fig. 3. Temperatures without cooling (thin lines) and optimal temperature set-points (thick lines) [°C].

distinguishing feature of our set-up, which required the introduction of a novel scheme, is that the electric energy price depends on the demand. This is the current trend in energy systems where pricing signals are used by grid operators to affect users consumption/generation. Further work is needed to extend our approach to a more realistic setting where demand is uncertain.

REFERENCES

Belluschi, F., Falsone, A., Ioli, D., Margellos, K., Garatti, S., and Prandini, M. (2016). Energy management for building district cooling: a distributed approach to resource sharing. *preprint arXiv:1610.06332*.

Bertsekas, D.P. (1999). *Nonlinear programming*. Athena scientific.

Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J. (2010). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. and Trends in Machine Learning*, 3(1), 1–122.

Chang, T.H., Alizadeh, M., and Scaglione, A. (2013). Real-time power balancing via decentralized coordinated home energy scheduling. *IEEE Transactions on Smart Grid*, 4(3), 1490–1504.

Deng, K., Sun, Y., Chakraborty, A., Lu, Y., Brouwer, J., and Mehta, P.G. (2013). Optimal scheduling of chiller plant with thermal energy storage using mixed integer linear programming. In *American Control Conference*. Washington, DC, USA.

Deori, L., Margellos, K., and Prandini, M. (2016). On decentralized convex optimization in a multi-agent setting with separable constraints and its application to optimal charging of electric vehicles. In *IEEE Conference on Decision and Control*. Las Vegas, USA.

Deori, L., Margellos, K., and Prandini, M. (20167). Regularized Jacobi iteration for decentralized convex optimization with separable constraints. *preprint arXiv:1604.07814*.

Gharesifard, B., Basar, T., and Dominguez-Garcia, A. (2016). Price-based coordinated aggregation of networked distributed energy resources. *IEEE Transactions on Automatic Control*, 61(10), 2936–2946.

Grammatico, S., Parise, F., Colombino, M., and Lygeros, J. (2016). Decentralized convergence to nash equilibria in constrained deterministic mean field control. *IEEE Transactions on Automatic Control*, 61(11), 3315–3329.

Ioli, D., Falsone, A., and Prandini, M. (2015a). An iterative scheme to hierarchically structured optimal energy management of a microgrid. In *IEEE Conference on Decision and Control*. Osaka, Japan.

Ioli, D., Falsone, A., and Prandini, M. (2015b). Optimal energy management of a building cooling system with thermal storage: A convex formulation. In *IFAC Symposium on Advanced Control of Chemical Processes*. Whistler, Canada.

Kranning, M., Chu, E., Lavaei, J., and Boyd, S. (2014). Dynamic Network Energy Management via Proximal Message Passing. *Foundations and Trends in Optimization*, 1(2), 70–122.

Ma, Y., Kelman, A., Daly, A., and Borrelli, F. (2012). Predictive control for energy efficient buildings with thermal storage: Modeling, simulation and experiments. *IEEE Control Systems Magazine*, 1(1), 45–66.

Ma, Z., Callaway, D., and Hiskens, I. (2013). Decentralized charging control of large populations of plug-in electric vehicles. *IEEE Transactions on Control Systems Technology*, 21(1), 67–78.

Pérez-Lombard, L., Ortiz, J., and Pout, C. (2008). A review on buildings energy consumption information. *Energy and Buildings*, 40(3), 394–398.

Scherer, H., Pasamontes, M., Guzmán, J., Álvarez, J., Camponogara, E., and Normey-Rico, J. (2014). Efficient building energy management using distributed model predictive control. *Journal of Process Control*, 24(6), 740 – 749.

Siroky, J., Oldewurtel, F., Cigler, J., and Privara, S. (2011). Experimental analysis of model predictive control for an energy efficient building heating system. *Applied Energy*, 88(9), 3079–3087.

Telaretti, E., Ippolito, M.G., Dusonchet, L., Favuzza, S., Sanseverino, E.R., Massaro, F., Zizzo, G., and Pecoraro, G. (2015). New dispatching strategy for the integration of active-demand and distributed storage in the electricity market. In *AEIT International Annual Conference*.

Teleke, S., Baran, M.E., Bhattacharya, S., and Huang, A.Q. (2010). Optimal control of battery energy storage for wind farm dispatching. *IEEE Transactions on Energy Conversion*, 25(3), 787–794.

Vignali, R.M., Borghesan, F., Piroddi, L., Strelec, M., and Prandini, M. (2017). Energy management of a building cooling system with thermal storage: An approximate dynamic programming solution. *IEEE Transactions on Automation Science and Engineering*, 14(2), 619–633.

Wang, Z., Gu, C., Li, F., Bale, P., and Sun, H. (2013). Active demand response using shared energy storage for household energy management. *IEEE Transactions on Smart Grid*, 4(4), 1888–1897.