

A decentralized approach to multi-agent MILPs: finite-time feasibility and performance guarantees

Alessandro Falsone^a, Kostas Margellos^b, Maria Prandini^a

^a*Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Via Ponzio 34/5, 20133 Milano, Italy*

^b*Department of Engineering Science, University of Oxford, Parks Road, Oxford, OX1 3PJ, United Kingdom*

Abstract

We address the optimization of a large scale multi-agent system where each agent has discrete and/or continuous decision variables that need to be set so as to optimize the sum of linear local cost functions, in presence of linear local and global constraints. The problem reduces to a Mixed Integer Linear Program (MILP) that is here addressed according to a decentralized iterative scheme based on dual decomposition, where each agent determines its decision vector by solving a smaller MILP involving its local cost function and constraint given some dual variable, whereas a central unit enforces the global coupling constraint by updating the dual variable based on the tentative primal solutions of all agents. An appropriate tightening of the coupling constraint through iterations allows to obtain a solution that is feasible for the original MILP. The proposed approach is inspired by a recent paper to the MILP approximate solution via dual decomposition and constraint tightening, but shows finite-time convergence to a feasible solution and provides sharper performance guarantees by means of an adaptive tightening. The two approaches are compared on a plug-in electric vehicles optimal charging problem.

Key words: MILP, decentralized optimization, multi-agent networks, electric vehicles.

1 Introduction

In this paper we are concerned with the optimization of a large-scale system composed of multiple agents, each one characterized by its set of decision variables that should be chosen so as to solve a constrained optimization problem where the agents' decisions are coupled by some global constraint. More specifically, the goal is to minimize the sum of local linear cost functions, subject to local polyhedral constraints and a global linear constraint. As in the inspiring work [18], we consider a framework where decision variables can have both continuous and discrete components, a feature that makes the problem challenging.

Let m denote the number of agents. Then, the optimization

* Research was supported by the European Commission under the project UnCoVerCPS, grant number 643921, and by EPSRC UK under the grant EP/P03277X/1.

**Corresponding author A. Falsone. Tel. +39-02-23994028. Fax +39-02-23993412.

Email addresses:

alessandro.falsone@polimi.it (Alessandro Falsone),
kostas.margellos@eng.ox.ac.uk (Kostas Margellos),
maria.prandini@polimi.it (Maria Prandini).

problem takes the form of the following Mixed Integer Linear Program (MILP):

$$\begin{aligned} \min_{x_1, \dots, x_m} \quad & \sum_{i=1}^m c_i^\top x_i \\ \text{subject to:} \quad & \sum_{i=1}^m A_i x_i \leq b \\ & x_i \in X_i, \quad i = 1, \dots, m \end{aligned} \quad (\mathcal{P})$$

where, for all $i = 1, \dots, m$, $x_i \in \mathbb{R}^{n_i}$ is the decision vector of agent i , $c_i^\top x_i$ its local cost, and $X_i = \{x_i \in \mathbb{R}^{n_{c,i}} \times \mathbb{Z}^{n_{d,i}} : D_i x_i \leq d_i\}$ its local constraint set defined by a matrix D_i and a vector d_i of appropriate dimensions, $n_{c,i}$ being the number of continuous decision variables and $n_{d,i}$ the number of discrete ones, with $n_{c,i} + n_{d,i} = n_i$. The coupling constraint $\sum_{i=1}^m A_i x_i \leq b$ is defined by matrices $A_i \in \mathbb{R}^p \times \mathbb{R}^{n_i}$, $i = 1, \dots, m$, and a p -dimensional vector $b \in \mathbb{R}^p$. Note that all inequalities involving vectors have to be intended component-wise.

Despite the advances in numerical methods for integer optimization, when the number of agents is large, the presence of discrete decision variables makes the optimization problem hard to solve, and calls for some de-

composition into lower scale MILPs, as suggested in [18].

A common practice to handle problems of the form of \mathcal{P} consists in first dualizing the coupling constraint introducing a vector $\lambda \in \mathbb{R}^p$ of p Lagrange multipliers and solving the dual program

$$\max_{\lambda \geq 0} -\lambda^\top b + \sum_{i=1}^m \min_{x_i \in X_i} (c_i^\top + \lambda^\top A_i)x_i, \quad (\mathcal{D})$$

to obtain λ^* , and then constructing a primal solution $x(\lambda^*) = [x_1(\lambda^*)^\top \cdots x_m(\lambda^*)^\top]^\top$ by solving m MILPs given by:

$$x_i(\lambda) \in \arg \min_{x_i \in \text{vert}(X_i)} (c_i^\top + \lambda^\top A_i)x_i, \quad (1)$$

where the search within the closed constraint polyhedral set X_i can be confined to its set of vertices $\text{vert}(X_i)$ since the cost function is linear. Unfortunately, while this procedure guarantees $x(\lambda^*)$ to satisfy the local constraints since $x_i(\lambda^*) \in X_i$ for all $i = 1, \dots, m$, it does not guarantee the satisfaction of the coupling constraint. An example illustrating this case is reported in Part 1 of Example 1 in Appendix A.

A way to enforce the satisfaction of the coupling constraint is to follow the approach in [16], where the dual program \mathcal{D} is solved via a particular iterative methodology, namely, the subgradient algorithm. At each iteration of the subgradient algorithm a tentative primal solution is generated by every agent. By appropriately averaging the tentative solutions across iterations (see [16, pag. 117]), one can obtain a solution that satisfies the joint constraint. However, when discrete decision variables are present, such solution does not necessarily satisfy also the local constraints. Specifically, letting $\text{conv}(X_i)$ denote the convex hull of X_i , $i = 1, \dots, m$, if we apply the above procedure to \mathcal{P} , we obtain an optimal solution x_{LP}^* to the following Linear Program (LP):

$$\begin{aligned} \min_{x_1, \dots, x_m} \quad & \sum_{i=1}^m c_i^\top x_i \\ \text{subject to:} \quad & \sum_{i=1}^m A_i x_i \leq b \\ & x_i \in \text{conv}(X_i), \quad i = 1, \dots, m. \end{aligned} \quad (\mathcal{P}_{\text{LP}})$$

This fact is true because the dual of the convexified problem \mathcal{P}_{LP} coincides with the dual of \mathcal{P} and is given by \mathcal{D} (see [11] for a proof). Clearly $x_{\text{LP}}^* \in \text{conv}(X_1) \times \cdots \times \text{conv}(X_m)$ does not necessarily imply that $x_{\text{LP}}^* \in X_1 \times X_2 \times \cdots \times X_m$. Therefore the solution x_{LP}^* recovered using [16] satisfies the coupling constraint but not necessarily the local constraints. An alternative approach for finding an optimal solution to the primal-dual pair $\mathcal{P}_{\text{LP}}\text{-}\mathcal{D}$ is to exploit the column generation algorithm

(see [13]). Even in this case however the procedure converges to a solution x_{LP}^* of \mathcal{P}_{LP} , which is not guaranteed to be feasible for the local constraints in \mathcal{P} . An example in which the solution to \mathcal{P}_{LP} is not feasible for the local constraints is reported in Part 2 of Example 1 in Appendix A.

For these reasons recovery procedures for MILPs are usually composed of two steps: a tentative solution that is not feasible for either the joint constraint or the local ones is first obtained exploiting one of the two procedures described above, and then a problem-specific heuristic is applied to recover a feasible solution for \mathcal{P} , see, e.g., [4,15].

Problems in the form of \mathcal{P} arise in different contexts like power plants generation scheduling [19] where the agents are the generation units with their on/off state modeled with binary variables and the joint constraint consists in energy balance equations, or buildings energy management [12], where the cost function is a cost related to power consumption and constraints are related to capacity, comfort, and actuation limits of each building. Other problems that fit the structure of \mathcal{P} are supply chain management [8], portfolio optimization for small investors [2], and plug-in electric vehicles [18]. In all these cases it is of major interest to guarantee that the derived (primal) solution is implementable in practice, which means that it must be feasible for \mathcal{P} .

Interestingly, a large class of dynamical systems involving both continuous and logic components can be modeled as a Mixed Logical Dynamical (MLD) system, using the terminology established in [3], which are described by linear equations and inequalities involving both discrete and continuous inputs and state variables. Finite horizon control for multiple MLD systems modeling interacting agents that are jointly optimizing a linear objective function while sharing some resources could be formulated as problem \mathcal{P} . Designing an iterative decentralized algorithm that is guaranteed to solve \mathcal{P} in finite time is then important for the development of decentralized model predictive control schemes for multi-agent MLD systems, since problem \mathcal{P} has to be repeatedly solved online within some time interval in that context.

Finite-time convergence to a solution which is at least feasible for \mathcal{P} is a desirable feature for most of the aforementioned applications. The main goal of this paper is to provide such a guarantee, which has up to now proven to be elusive.

1.1 Background

Problems in the form of \mathcal{P} have been investigated in [1], where the authors studied the behavior of the duality gap (i.e., the difference between the optimal value of \mathcal{P}

and \mathcal{D}) showing that it decreases relatively to the optimal value of \mathcal{P} as the number of agents grows. The same behavior has been observed in [4]. In the recent paper [18], the authors explored the connection between the solutions x_{LP}^* to the linear program \mathcal{P}_{LP} and $x(\lambda^*)$ recovered via (1) from the solution λ^* to the dual program \mathcal{D} . They proposed a method to recover a primal solution which is feasible for \mathcal{P} by using the dual optimal solution of a modified primal problem, obtained by tightening the coupling constraint by an appropriate amount.

We now recall those parts of [18] that are relevant for the developments in this paper.

Let $\rho \in \mathbb{R}^p$ with $\rho \geq 0$ and consider the following pair of primal-dual problems:

$$\begin{aligned} \min_{x_1, \dots, x_m} \quad & \sum_{i=1}^m c_i^\top x_i & (\mathcal{P}_{\text{LP}, \rho}) \\ \text{subject to:} \quad & \sum_{i=1}^m A_i x_i \leq b - \rho \\ & x_i \in \text{conv}(X_i), \quad i = 1, \dots, m \end{aligned}$$

and

$$\max_{\lambda \geq 0} \quad -\lambda^\top (b - \rho) + \sum_{i=1}^m \min_{x_i \in X_i} (c_i^\top + \lambda^\top A_i) x_i. \quad (\mathcal{D}_\rho)$$

$\mathcal{P}_{\text{LP}, \rho}$ constitutes a tightened version of \mathcal{P}_{LP} , whereas \mathcal{D}_ρ is the corresponding dual. For all $j = 1, \dots, p$, let $\tilde{\rho} \in \mathbb{R}^p$ be defined as follows:

$$[\tilde{\rho}]_j = p \max_{i \in \{1, \dots, m\}} \left\{ \max_{x_i \in X_i} [A_i]_j x_i - \min_{x_i \in X_i} [A_i]_j x_i \right\}, \quad (2)$$

where $[A_i]_j$ denotes the j -th row of A_i and $[\tilde{\rho}]_j$ the j -th entry of $\tilde{\rho}$.

Define $\tilde{\mathcal{P}}_{\text{LP}}$ and $\tilde{\mathcal{D}}$ as the primal-dual pair of optimization problems that are given by setting ρ equal to $\tilde{\rho}$ in $\mathcal{P}_{\text{LP}, \rho}$ and \mathcal{D}_ρ .

Assumption 1 (Existence and uniqueness, [18])
Problems $\tilde{\mathcal{P}}_{\text{LP}}$ and $\tilde{\mathcal{D}}$ have unique solutions $x_{\text{LP}, \tilde{\rho}}^*$ and $\lambda_{\tilde{\rho}}^*$.

Proposition 1 (Theorem 3.1 in [18]) *Let $\lambda_{\tilde{\rho}}^*$ be the solution to $\tilde{\mathcal{D}}$. Under Assumption 1, we have that any $x(\lambda_{\tilde{\rho}}^*)$ satisfying (1), is feasible for \mathcal{P} .*

The proof of Proposition 1 rests on Theorem 2.5 in [18]. Example 2.6 in [18] shows how Theorem 2.5 in [18], and therefore also Proposition 1, strongly depend on the uniqueness part of Assumption 1. Note, however, that in case $\tilde{\mathcal{P}}_{\text{LP}}$ has multiple solutions, then a small perturbation in its cost coefficients will render its solution unique,

thus making Assumption 1 fulfilled again. We refer the reader to [18] for further details.

Let us define

$$\tilde{\gamma} = p \max_{i \in \{1, \dots, m\}} \left\{ \max_{x_i \in X_i} c_i^\top x_i - \min_{x_i \in X_i} c_i^\top x_i \right\}. \quad (3)$$

Consider the following assumption:

Assumption 2 (Slater, [18]) *There exist a scalar $\zeta > 0$ and $\hat{x}_i \in \text{conv}(X_i)$ for all $i = 1, \dots, m$, such that $\sum_{i=1}^m A_i \hat{x}_i \leq b - \tilde{\rho} - m\zeta \mathbb{1}$, where $\mathbb{1} \in \mathbb{R}^p$ is a vector whose elements are equal to one.*

Then, the sub-optimality level of the approximate solution $x(\lambda_{\tilde{\rho}}^*)$ to \mathcal{P} can be quantified as follows:

Proposition 2 (Theorem 3.3 in [18]) *Let $\lambda_{\tilde{\rho}}^*$ be the solution to $\tilde{\mathcal{D}}$. Under Assumptions 1 and 2, we have that $x(\lambda_{\tilde{\rho}}^*)$ derived from (1) with $\lambda = \lambda_{\tilde{\rho}}^*$ satisfies*

$$\sum_{i=1}^m c_i^\top x_i(\lambda_{\tilde{\rho}}^*) - J_{\mathcal{P}}^* \leq \tilde{\gamma} + \frac{\|\tilde{\rho}\|_\infty}{p\zeta} \tilde{\gamma}, \quad (4)$$

where $J_{\mathcal{P}}^*$ is the optimal cost of \mathcal{P} .

Assumption 2 is used to estimate (through Lemma 1 in [14]) the norm of the optimal solution $\lambda_{\tilde{\rho}}^*$ to the dual problem, which is needed in the proof of Proposition 2 to derive the performance bound. Assumption 2 is instead not needed to prove feasibility.

Note that both Proposition 1 on feasibility and Proposition 2 on optimality require knowledge of the dual solution $\lambda_{\tilde{\rho}}^*$. This may pose some issues if $\lambda_{\tilde{\rho}}^*$ cannot be computed centrally, which is the case, e.g., when the agents are not willing to share with some central entity their private information coded in their local cost and constraint set. In those cases, the value of $\lambda_{\tilde{\rho}}^*$ can only be achieved asymptotically using a decentralized/distributed scheme to solve $\tilde{\mathcal{D}}$.

1.2 Contribution of this paper

In this paper, we propose a decentralized iterative procedure which computes in a finite number of iterations a solution that is feasible for the optimization problem \mathcal{P} . We also provide performance guarantees quantifying the sub-optimality level of our solution with respect to the optimal one of \mathcal{P} .

The proposed iterative method is inspired by the work in [18]. As in [18], we exploit some tightening of the coupling constraint to enforce feasibility. The amount of

tightening introduced in our method is adaptively chosen throughout the iterations, based on the explored candidate solutions $x_i \in X_i$, $i = 1, \dots, m$, and is guaranteed to be lower than or equal to the worst-case tightening $\tilde{\rho}$ adopted in [18] which is obtained by letting x_i vary over the whole set X_i (see equation (2)). Note that a large value of $\tilde{\rho}$ may prevent $\tilde{\mathcal{P}}_{\text{LP}}$ to be feasible thus hampering the applicability of the approach in [18]. This is easy to understand if b in the coupling constraint is interpreted as the maximum available amount of some (shared) resource: if such an amount is reduced by ρ and ρ is large, then, it might be that the remaining amount of resource $b - \rho$ is not enough to satisfy the local constraints of the agents, thus resulting in infeasibility. A less conservative way of selecting the amount of tightening as in our method may preserve the feasibility in the tightened problem, thus making our approach applicable to cases where the approach in [18] is not. This is shown in Section 5 where the plug-in electric vehicles charging problem originally presented in [18] is considered as a case study: in the vehicle to grid setup with $m = 250$ vehicles, when the maximum power b that the network can deliver is reduced, then, $\tilde{\mathcal{P}}_{\text{LP}}$ becomes infeasible and, hence, the approach in [18] cannot be applied, whereas our method remains applicable because it introduces a smaller tightening. We can then claim that our method can be applied to a larger class of problems than the method in [18]. Furthermore, when both methods can be applied but the tightening of our method is smaller, performance guarantees are better for our solution as quantified through the bound on the obtained improvement derived in Remark 1 at the end of Section 2. This is also demonstrated in Section 5: in all the 1000 instances of the plug-in electric vehicles charging problem generated though some random perturbation of the involved parameters, tightening is smaller and performance is better in our method.

Finite-time convergence is certainly the main feature of our approach, which makes it attractive for various applications and for MPC in particular. Additionally, finite-time convergence has a direct impact on computational complexity, which is alleviated with respect to the approach in [18]. This is clarified next.

Both methods exploit the structural properties of the MILP \mathcal{P} to cope with its combinatorial complexity by decomposing it into m smaller MILPs with fewer discrete decision variables. However, from the point of view of the resolution of the dual program, a different computational complexity arises in the two methods. In the simulation section of [18] a subgradient algorithm that asymptotically converges to the dual optimal solution is employed. Therefore, the approach in [18] would need, in principle, to solve the m MILPs for an infinite number of iterations, whilst we only have to solve the m MILPs a finite number of times. This clearly shows that the computational complexity needed for our method to solve \mathcal{P} is lower compared to that in [18].

Algorithm 1 Decentralized MILP

```

1:  $\lambda(0) = 0$ 
2:  $\bar{s}_i(0) = -\infty$ ,  $i = 1, \dots, m$ 
3:  $\underline{s}_i(0) = +\infty$ ,  $i = 1, \dots, m$ 
4:  $k = 0$ 
5: repeat
6:   for  $i = 1$  to  $m$  do
7:      $x_i(k+1) \leftarrow \arg \min_{x_i \in \text{vert}(X_i)} (c_i^\top + \lambda(k)^\top A_i) x_i$ 
8:   end for
9:    $\bar{s}_i(k+1) = \max\{\bar{s}_i(k), A_i x_i(k+1)\}$ ,  $i = 1, \dots, m$ 
10:   $\underline{s}_i(k+1) = \min\{\underline{s}_i(k), A_i x_i(k+1)\}$ ,  $i = 1, \dots, m$ 
11:   $\rho_i(k+1) = \bar{s}_i(k+1) - \underline{s}_i(k+1)$ ,  $i = 1, \dots, m$ 
12:   $\rho(k+1) = p \max\{\rho_1(k+1), \dots, \rho_m(k+1)\}$ 
13:   $\lambda(k+1) = \left[ \lambda(k) + \alpha(k) \left( \sum_{i=1}^m A_i x_i(k+1) - b + \rho(k+1) \right) \right]_+$ 
14:   $k \leftarrow k + 1$ 
15: until some stopping criterion is met.

```

In summary, the differences between the approach proposed here and the one presented in [18] are:

- (1) adaptive versus worst-case constraint tightening, with implications in terms of applicability to a larger class of problems and better performance guarantees when both approaches are applicable;
- (2) finite-time versus asymptotic guarantees, with implications in terms of computational complexity.

Notably, both methods allow agents to preserve privacy of their local information, since they do not have to share either their cost coefficients or their local constraints.

2 Proposed approach

We next introduce Algorithm 1 for decentralized computation in a finite number of iterations of an approximate solution to \mathcal{P} that is feasible and improves over the solution in [18] both in terms of amount of tightening and performance guarantees.

Algorithm 1 is a variant of the dual subgradient algorithm. As the standard dual subgradient method, it includes two main steps: step 7 in which a subgradient of the dual objective function is computed by fixing the dual variables and minimizing the Lagrangian with respect to the primal variables, and step 13 which involves a dual update step with step size equal to $\alpha(k)$, and a projection onto the non-negative orthant (in Algorithm 1 $[\cdot]_+$ denotes the projection operator onto the p -dimensional non-negative orthant \mathbb{R}_+^p). The operators \max and \min appearing in steps 9, 10, and 12 of Algorithm 1 with arguments in \mathbb{R}^p are meant to be applied component-wise. The sequence $\{\alpha(k)\}_{k \geq 0}$ is chosen so as to satisfy $\lim_{k \rightarrow \infty} \alpha(k) = 0$ and $\sum_{k=0}^{\infty} \alpha(k) = \infty$, as requested in the standard dual subgradient method to

achieve asymptotic convergence. Furthermore, in order to guarantee that the solution to step 7 of Algorithm 1 is well-defined, we impose the following assumption on \mathcal{P} :

Assumption 3 (Boundedness) *The polyhedral sets X_i , $i = 1, \dots, m$, in problem \mathcal{P} are bounded.*

If $\arg \min_{x_i \in \text{vert}(X_i)} (c_i^\top + \lambda(k)^\top A_i) x_i$ in step 7 is a set of cardinality larger than 1, then, a deterministic tie-break rule is applied to choose a value for $x_i(k+1)$.

Algorithm 1 is conceived to be implemented in a decentralized scheme where, at each iteration k , every agent i updates its local tentative solution $x_i(k+1)$ and communicates $A_i x_i(k+1)$ to some central unit that is in charge of the update of the dual variable. The tentative value $\lambda(k+1)$ for the dual variable is then broadcast to all agents. Note that agents do not need to communicate to the central unit their private information regarding their local constraint set and cost but only their tentative solution $x_i(k)$.

The tentative primal solutions $x_i(k+1)$, $i = 1, \dots, m$, computed at step 7 are used in Algorithm 1 by the central unit to determine the amount of tightening $\rho(k+1)$ entering step 13. The value of $\rho(k+1)$ is progressively refined through iterations based only on those values of $x_i \in X_i$, $i = 1, \dots, m$, that are actually considered as candidate primal solutions, and not based on the whole sets X_i , $i = 1, \dots, m$. This reduces conservativeness in the amount of tightening and also in the performance bound of the feasible, yet suboptimal, primal solution.

A further reduction in the level of conservativeness can be achieved by assigning to $[\rho_i(k+1)]_j$ in step 12 of Algorithm 1 the (less conservative) sum of the p -largest $[\rho_i(k+1)]_j$, for all $j = 1, \dots, p$. Further discussion is provided after the proof of Proposition 3. Although this is not discussed in [18], also the results in [18] can be modified to use this less conservative bound. For a fair comparison, this modification is not included in Algorithm 1 as well.

Algorithm 1 terminates after a given stopping criteria is met at the level of the central unit, e.g., if for a given number of subsequent iterations $x(k) = [x_1(k)^\top \dots x_m(k)^\top]^\top$ satisfies the coupling constraint. As shown in the numerical study in Section 5, variants of Algorithm 1 can be conceived to get an improved solution in the same number of iterations of Algorithm 1. The agents should however share with the central entity additional information on their local cost, thus partly compromising privacy preservation.

As for the initialization of Algorithm 1, $\lambda(0)$ is set equal to 0 so that at iteration $k = 0$ each agent i computes its locally optimal solution

$$x_i(1) \leftarrow \arg \min_{x_i \in \text{vert}(X_i)} c_i^\top x_i.$$

Since $\rho(1) = 0$, if the local solutions $x_i(1)$, $i = 1, \dots, m$, satisfy the coupling constraint (and they hence are optimal for the original problem \mathcal{P}), then, Algorithm 1 will terminate since λ will remain 0, and the agents will stick to their locally optimal solutions.

Before stating the feasibility and performance guarantees of the solution computed by Algorithm 1, we need to introduce some further quantities and assumptions.

Let us define for any $k \geq 1$

$$\gamma(k) = p \max_{i \in \{1, \dots, m\}} \left\{ \max_{r \leq k} c_i^\top x_i(r) - \min_{r \leq k} c_i^\top x_i(r) \right\}, \quad (5)$$

where $\{x_i(r)\}_{r \geq 1}$, $i = 1, \dots, m$, are the tentative primal solutions computed at step 7.

Due to Assumption 3, for any $i = 1, \dots, m$, $\text{conv}(X_i)$ is a bounded polyhedron. If it is also non-empty, then $\text{vert}(X_i)$ is a non-empty finite set (see Corollaries 2.1 and 2.2 together with Theorem 2.3 in [6, Chapter 2]). As a consequence, the sequence $\{\gamma(k)\}_{k \geq 1}$ takes values in a finite set. Since this is a monotonically non-decreasing sequence, it converges in finite-time to some value $\bar{\gamma}$. The same reasoning can be applied to show that the sequence $\{\rho(k)\}_{k \geq 1}$, iteratively computed in Algorithm 1 (see step 12), and given by

$$[\rho(k)]_j = p \max_{i \in \{1, \dots, m\}} \left\{ \max_{r \leq k} [A_i]_j x_i(r) - \min_{r \leq k} [A_i]_j x_i(r) \right\},$$

for $j = 1, \dots, p$, converges in finite-time to some $\bar{\rho}$ since it takes values in a finite set and is (component-wise) monotonically non-decreasing. Note that the limiting values $\bar{\rho}$ and $\bar{\gamma}$ for $\{\rho(k)\}_{k \geq 1}$ and $\{\gamma(k)\}_{k \geq 1}$ satisfy $\bar{\rho} \leq \bar{\rho}$ and $\bar{\gamma} \leq \bar{\gamma}$ where $\bar{\rho}$ and $\bar{\gamma}$ are defined in (2) and (3).

Define $\bar{\mathcal{P}}_{\text{LP}}$ and $\bar{\mathcal{D}}$ as the primal-dual pair of optimization problems that are given by setting ρ equal to $\bar{\rho}$ in $\mathcal{P}_{\text{LP}, \rho}$ and \mathcal{D}_ρ .

In order to state the feasibility and performance properties of Algorithm 1, besides Assumption 3, the following two further assumptions are needed.

Assumption 4 (Existence and uniqueness)

Problems $\bar{\mathcal{P}}_{\text{LP}}$ and $\bar{\mathcal{D}}$ have unique solutions \bar{x}_{LP}^ and $\bar{\lambda}^*$.*

Assumption 5 (Slater) *There exists a scalar $\zeta > 0$ and $\hat{x}_i \in \text{conv}(X_i)$ for all $i = 1, \dots, m$, such that $\sum_{i=1}^m A_i \hat{x}_i \leq b - \bar{\rho} - m\zeta \mathbf{1}$.*

Note that Assumptions 4 and 5 are similar to Assumptions 1 and 2, respectively. However, owing to the fact that $\bar{\rho} \leq \bar{\rho}$, imposing Assumptions 4 and 5 in place of Assumptions 1 and 2 makes Algorithm 1 applicable to a

larger class of problems with respect to the approach in [18].

The discussion about the necessity and plausibility of these assumptions follows closely that related to Assumptions 1 and 2 in Section 1.1 and is here omitted.

We are now in a position to state the two main results of the paper.

Theorem 3 (Finite-time feasibility) *Under Assumptions 3 and 4, there exists a finite iteration index K such that, for all $k \geq K$, $x(k) = [x_1(k)^\top \dots x_m(k)^\top]^\top$, where $x_i(k)$, $i = 1, \dots, m$, are computed by Algorithm 1, is a feasible solution for problem \mathcal{P} , i.e., $\sum_{i=1}^m A_i x_i(k) \leq b$, $k \geq K$ and $x_i(k) \in X_i$, $i = 1, \dots, m$.*

Theorem 4 (Performance guarantees) *Under Assumptions 3-5, there exists a finite iteration index K such that, for all $k \geq K$, $x(k) = [x_1(k)^\top \dots x_m(k)^\top]^\top$, where $x_i(k)$, $i = 1, \dots, m$, are computed by Algorithm 1, is a feasible solution for problem \mathcal{P} that satisfies the following performance bound:*

$$\sum_{i=1}^m c_i^\top x_i(k) - J_{\mathcal{P}}^* \leq \bar{\gamma} + \frac{\|\bar{\rho}\|_\infty}{p\zeta} \tilde{\gamma}. \quad (6)$$

By a direct comparison of (4) and (6) we can see that the bound in (6) is no worse than (4) due to the fact that $\bar{\rho} \leq \bar{\rho}$ and $\bar{\gamma} \leq \tilde{\gamma}$.

In the following remark, performance improvement is quantified when both methods are applicable and $\bar{\rho} < \bar{\rho}$.

Remark 1 (Performance improvement versus [18]) *Suppose that Assumption 1 (and, hence Assumption 4) is satisfied. Let $\bar{\rho} < \bar{\rho}$.*

Consider $\hat{x}_i \in \text{conv}(X_i)$, $i = 1, \dots, m$, such that Assumption 2 is satisfied with a given $\tilde{\zeta}$. Then, Assumption 5 is satisfied with the same \hat{x}_i , $i = 1, \dots, m$, and $\tilde{\zeta} = \tilde{\zeta} + \frac{1}{m} \min_{j=1, \dots, p} \{[\bar{\rho}]_j - [\bar{\rho}]_j\} > \tilde{\zeta}$. This implies that our performance bound in (6) is tighter than the one in (4) by an amount equal to

$$\tilde{\gamma} - \bar{\gamma} + \frac{\tilde{\gamma}}{p} \left[\frac{\|\bar{\rho}\|_\infty}{\tilde{\zeta}} - \frac{\|\bar{\rho}\|_\infty}{\tilde{\zeta}} \right] > 0,$$

where we used the fact that $\bar{\rho} < \bar{\rho}$, $\tilde{\zeta} > \tilde{\zeta}$ and $\bar{\gamma} \leq \tilde{\gamma}$.

Note that if $\bar{\mathcal{P}}_{\text{LP}}$ is not feasible for the resulting $\bar{\rho}$, then $\sum_{i=1}^m A_i x_i(k+1) - b + \rho(k+1)$ is bounded below by some positive constant for a sufficiently high k given that $\rho(k+1)$ converges to $\bar{\rho}$. Since $\sum_{k=0}^\infty \alpha(k) = \infty$, step 13 of Algorithm 1 will then produce a $\{\lambda(k)\}_{k \geq 0}$ sequence diverging towards $+\infty$. Therefore, observing a

component of $\lambda(k)$ which diverges as k increases is an indication that the existence part of Assumption 4 is not satisfied.

3 Proof of the main results

3.1 Preliminary results

Proposition 5 (Dual asymptotic convergence)

Under Assumptions 3 and 4, the Lagrange multiplier sequence $\{\lambda(k)\}_{k \geq 0}$ generated by Algorithm 1 converges to an optimal solution of $\bar{\mathcal{D}}$.

Proof. As discussed after equation (5), there exists a $K \in \mathbb{N}$ such that for all $k \geq K$ we have that the tightening coefficient $\rho(k)$ computed in Algorithm 1 becomes constant and equal to $\bar{\rho}$. Therefore, for any $k \geq K$, Algorithm 1 reduces to the following two steps

$$\begin{aligned} x_i(k+1) &\in \arg \min_{x_i \in \text{vert}(X_i)} (c_i^\top + \lambda(k)^\top A_i) x_i \\ \lambda(k+1) &= \left[\lambda(k) + \alpha(k) \left(\sum_{i=1}^m A_i x_i(k+1) - b + \bar{\rho} \right) \right]_+^{(8)} \end{aligned} \quad (7)$$

which constitute a gradient ascent iteration for $\bar{\mathcal{D}}$. According to [5], the sequence $\{\lambda(k)\}_{k \geq 0}$ generated by the iterative procedure (7)-(8) is guaranteed to converge to the (unique under Assumption 4) optimal solution of $\bar{\mathcal{D}}$.

Note that this result requires only uniqueness of the optimal solution of $\bar{\mathcal{D}}$. Uniqueness of the optimal solution to $\bar{\mathcal{P}}_{\text{LP}}$ is not necessary. \square

Lemma 1 (Robustness against cost perturbation)

Let P be a non-empty bounded polyhedron. Consider the linear program $\min_{x \in P} (c^\top + \delta^\top) x$, where δ is a perturbation in the cost coefficients. Define the set of optimal solutions as $\mathcal{X}(\delta)$. There always exists an $\varepsilon > 0$ such that for all δ satisfying $\|\delta\| < \varepsilon$, we have $\mathcal{X}(\delta) \subseteq \mathcal{X}(0)$.

Proof. Let $u(\delta) = \min_{x \in P} (c^\top + \delta^\top) x$. Since P is a bounded polyhedron, the minimum is always attained and $u(\delta)$ is finite for any value of δ . The set $\mathcal{X}(\delta)$ can be defined as

$$\mathcal{X}(\delta) = \{x \in P : (c^\top + \delta^\top) x \leq u(\delta)\}, \quad (9)$$

which is a non-empty polyhedron. As such, it can be described as the convex hull of its vertices (see Theorem 2.9 in [6, Chapter 2]), which are also vertices of P (Theorem 2.7 in [6, Chapter 2]).

Let $V = \text{vert}(P)$ and $V_\delta = \text{vert}(\mathcal{X}(\delta)) \subseteq V$. Consider $\delta = 0$.

If $V_0 = V$, then, given the fact that, for any δ , $\mathcal{X}(\delta)$ is the convex hull of V_δ and $V_\delta \subseteq V = V_0$, we have trivially that $\mathcal{X}(\delta) \subseteq \mathcal{X}(0)$, for any δ .

Suppose now that $V_0 \subset V$. For any choice of $x^* \in V_0$ and $x \in V \setminus V_0$, we have that $c^\top x^* < c^\top x$, or equivalently $c^\top(x^* - x) < 0$. Pick

$$\varepsilon = \min_{\substack{x^* \in V_0 \\ x \in V \setminus V_0}} -\frac{c^\top(x^* - x)}{\|x^* - x\|} \quad (10)$$

and let (\bar{x}^*, \bar{x}) be the corresponding minimizer. By construction, (10) is well defined since \bar{x}^* is different from \bar{x} . Since $c^\top(x^* - x) < 0$ for any $x^* \in V_0$ and $x \in V \setminus V_0$, we have that $\varepsilon > 0$. Moreover, for any $x^* \in V_0$ and $x \in V \setminus V_0$, if δ satisfies $\|\delta\| < \varepsilon$, then

$$\begin{aligned} (c^\top + \delta^\top)(x^* - x) &= c^\top(x^* - x) + \delta^\top(x^* - x) \\ &\leq c^\top(x^* - x) + \|\delta\| \|x^* - x\| \\ &< c^\top(x^* - x) + \varepsilon \|x^* - x\| \\ &\leq c^\top(x^* - x) + \left(-\frac{c^\top(x^* - x)}{\|x^* - x\|}\right) \|x^* - x\| \\ &= c^\top(x^* - x) - c^\top(x^* - x) = 0, \end{aligned} \quad (11)$$

where the first inequality is given by the fact that $u^\top v \leq |u^\top v|$ together with the Cauchy-Schwarz inequality $|u^\top v| \leq \|u\| \|v\|$, the second inequality is due to δ satisfying $\|\delta\| < \varepsilon$, and the third inequality is given by the definition of ε in (10).

By (9) and the definition of $u(\delta)$, for any point x_δ in the set V_δ , we have that $(c^\top + \delta^\top)x_\delta \leq (c^\top + \delta^\top)x$, for all $x \in V$, and therefore $(c^\top + \delta^\top)x_\delta \leq (c^\top + \delta^\top)x^*$ for any $x^* \in V_0 \subset V$. By (11), whenever $\|\delta\| < \varepsilon$, we have that $(c^\top + \delta^\top)x^* < (c^\top + \delta^\top)x$ for any choice of $x^* \in V_0$ and $x \in V \setminus V_0$, therefore $(c^\top + \delta^\top)x_\delta < (c^\top + \delta^\top)x$ for any $x \in V \setminus V_0$. Since the inequality is strict, we have that $x_\delta \notin V \setminus V_0$, which implies $x_\delta \in V_0$. Since this holds for any $x_\delta \in V_\delta$, we have that $V_\delta \subseteq V_0$.

Finally, given the fact that, for any δ , $\mathcal{X}(\delta)$ is the convex hull of V_δ and $V_\delta \subseteq V_0$, we have $\mathcal{X}(\delta) \subseteq \mathcal{X}(0)$, thus concluding the proof. \square

Exploiting Lemma 1, we shall show next that each $\{x_i(k)\}_{k \geq 1}$ sequence, $i = 1, \dots, m$, converges in finite-time to some set. Note that, for the subsequent result, only uniqueness of the optimal solution of $\bar{\mathcal{D}}$ is required.

Proposition 6 (Primal finite-time set convergence)

Under Assumptions 3 and 4, there exists a finite K such that for all $i = 1, \dots, m$ the tentative primal solution $x_i(k)$ generated by Algorithm 1 satisfies

$$x_i(k) \in \arg \min_{x_i \in \text{vert}(X_i)} (c_i^\top + \bar{\lambda}^{*\top} A_i) x_i, \quad k \geq K, \quad (12)$$

where $\bar{\lambda}^*$ is the limit value of the Lagrange multiplier sequence $\{\lambda(k)\}_{k \geq 0}$.

Proof. Consider agent i , with $i \in \{1, \dots, m\}$. We can characterize the solution $x_i(k)$ in step 7 of Algorithm 1 by performing the minimization over $\text{conv}(X_i)$ instead of $\text{vert}(X_i)$ since the problem is linear and by enlarging the set $\text{vert}(X_i)$ to $\text{conv}(X_i)$ we still obtain all minimizers that belong to $\text{vert}(X_i)$. Adding and subtracting $\bar{\lambda}^{*\top} A_i x_i$ to the cost, we then obtain

$$x_i(k) \in \arg \min_{x_i \in \text{conv}(X_i)} (c_i^\top + \bar{\lambda}^{*\top} A_i + (\lambda(k-1) - \bar{\lambda}^*)^\top A_i) x_i. \quad (13)$$

Set $\delta_i(k-1)^\top = (\lambda(k-1) - \bar{\lambda}^*)^\top A_i$, and let $\mathcal{X}_i(\delta_i(k-1))$ be the set of minimizers of (13) as a function of $\delta_i(k-1)$. By Lemma 1, we know that there exists an $\varepsilon_i > 0$ such that if $\|\delta_i(k-1)\| < \varepsilon_i$, then $\mathcal{X}_i(\delta_i(k-1)) \subseteq \mathcal{X}_i(0)$.

Since, by Proposition 5, the sequence $\{\lambda(k)\}_{k \geq 0}$ generated by Algorithm 1 converges to $\bar{\lambda}^*$, by definition of limit, we know that there exists a K_i such that $\|\delta_i(k-1)\| = \|(\lambda(k-1) - \bar{\lambda}^*)^\top A_i\| < \varepsilon_i$ for all $k \geq K_i$. Therefore, for every $k \geq K = \max\{K_1, \dots, K_m\}$, we have that $x_i(k) \in \mathcal{X}_i(0) = \arg \min_{x_i \in \text{conv}(X_i)} (c_i^\top + \bar{\lambda}^{*\top} A_i) x_i$, $i = 1, \dots, m$. This property jointly with the fact that $x_i(k) \in \text{vert}(X_i)$, $i = 1, \dots, m$, leads to (12), thus concluding the proof. \square

3.2 Proof of Theorems 3 and 4

Before discussing the proofs of Theorem 3 and 4 we shall emphasize that Theorem 2.5 in [18] and Lemma 1 in [14] are key for the following derivations.

Proof of Theorem 3. Theorem 2.5 of [18] establishes a relation between the solution \bar{x}_{LP}^* of $\bar{\mathcal{P}}_{\text{LP}}$ and the one recovered in (1) from the optimal solution $\bar{\lambda}^*$ of the dual optimization problem $\bar{\mathcal{D}}$. Specifically, it states that there exists a set of indices $I \subseteq \{1, \dots, m\}$ of cardinality at least $m - p$, such that $[\bar{x}_{\text{LP}}^*]^{(i)} = x_i(\bar{\lambda}^*)$ for all $i \in I$, where $[\bar{x}_{\text{LP}}^*]^{(i)}$ is the subvector of \bar{x}_{LP}^* corresponding to the i -th agent. Therefore, following the proof of Theo-

rem 3.1 in [18], we have that

$$\begin{aligned}
& \sum_{i=1}^m A_i x_i(\bar{\lambda}^*) \\
&= \sum_{i \in I} A_i x_i(\bar{\lambda}^*) + \sum_{i \in I^c} A_i x_i(\bar{\lambda}^*) \\
&= \sum_{i \in I} A_i [\bar{x}_{\text{LP}}^*]^{(i)} + \sum_{i \in I^c} A_i x_i(\bar{\lambda}^*) \\
&= \sum_{i=1}^m A_i [\bar{x}_{\text{LP}}^*]^{(i)} + \sum_{i \in I^c} A_i \left(x_i(\bar{\lambda}^*) - [\bar{x}_{\text{LP}}^*]^{(i)} \right) \\
&\leq b - \bar{\rho} + p \max_{i=1, \dots, m} \{ A_i x_i(\bar{\lambda}^*) - A_i [\bar{x}_{\text{LP}}^*]^{(i)} \}, \quad (14)
\end{aligned}$$

where $I^c = \{1, \dots, m\} \setminus I$, and $b - \bar{\rho}$ constitutes an upper bound for $\sum_{i=1}^m A_i [\bar{x}_{\text{LP}}^*]^{(i)}$ given that \bar{x}_{LP}^* is feasible for $\bar{\mathcal{P}}_{\text{LP}}$.

According to [16, pag. 117], the component $[x_{\text{LP}}^*]^{(i)}$ of the (unique, under Assumption 4) solution \bar{x}_{LP}^* to $\bar{\mathcal{P}}_{\text{LP}}$ is the limit point of the sequence $\{\tilde{x}_i(k)\}_{k \geq 1}$, defined as

$$\tilde{x}_i(k) = \frac{\sum_{r=1}^{k-1} \alpha(r) x_i(r+1)}{\sum_{r=1}^{k-1} \alpha(r)}.$$

By linearity, for all $k \geq 0$, we have that

$$\begin{aligned}
A_i \tilde{x}_i(k) &= \frac{\sum_{r=1}^{k-1} \alpha(r) A_i x_i(r+1)}{\sum_{r=1}^{k-1} \alpha(r)} \\
&\geq \min_{r \leq k} A_i x_i(r) \\
&= \underline{s}_i(k) \\
&\geq \underline{s}_i,
\end{aligned}$$

where the first inequality is due to the fact that all $\alpha(k)$ are positive and the second equality follows from step 10 of Algorithm 1. In the final inequality, $\underline{s}_i(k)$ is lower bounded by \underline{s}_i , that denotes the limiting value of the non-increasing finite-valued sequence $\{\underline{s}_i(k)\}_{k \geq 0}$. Recall that all inequalities have to be intended component-wise. By taking the limit for $k \rightarrow \infty$, we also have that

$$A_i [\bar{x}_{\text{LP}}^*]^{(i)} \geq \underline{s}_i. \quad (15)$$

By Proposition 6, there exists a finite iteration index K such that $x_i(k)$ satisfies (12). Since (14) holds for any choice of $x_i(\bar{\lambda}^*)$ which minimizes $(c_i^\top + \bar{\lambda}^{*\top} A_i) x_i$ over $\text{vert}(X_i)$, if $k \geq K$, then we can choose $x_i(\bar{\lambda}^*) = x_i(k)$.

Therefore, for all $k \geq K$, (14) becomes

$$\begin{aligned}
& \sum_{i=1}^m A_i x_i(k) \\
&\leq b - \bar{\rho} + p \max_{i=1, \dots, m} \{ A_i x_i(k) - A_i [x_{\text{LP}}^*]^{(i)} \} \\
&\leq b - \bar{\rho} + p \max_{i=1, \dots, m} \left\{ \max_{r \leq k} A_i x_i(r) - A_i [x_{\text{LP}}^*]^{(i)} \right\} \\
&= b - \bar{\rho} + p \max_{i=1, \dots, m} \{ \bar{s}_i(k) - A_i [x_{\text{LP}}^*]^{(i)} \} \\
&\leq b - \bar{\rho} + p \max_{i=1, \dots, m} \{ \bar{s}_i - \underline{s}_i \} \\
&= b, \quad (16)
\end{aligned}$$

where the second inequality is obtained by taking the maximum up to k , the first equality is due to step 9 of Algorithm 1, the third inequality is due to the fact that \bar{s}_i is the limiting value of the non-decreasing finite-valued sequence $\{\bar{s}_i(k)\}_{k \geq 1}$ together with (15), and the last equality comes from the definition of $\rho(k) = p \max\{\rho_1(k), \dots, \rho_m(k)\}$ where $\rho_i(k) = \bar{s}_i(k) - \underline{s}_i(k)$.

From (16) we have that, for any $k \geq K$, the iterates $x_i(k)$, $i = 1, \dots, m$, generated by Algorithm 1 provide a feasible solution for \mathcal{P} , thus concluding the proof. \square

As mentioned in Section 2, we can make Algorithm 1 less conservative by assigning to $[\rho_i(k+1)]_j$ in step 12 of Algorithm 1 the sum of the p -largest $[\rho_i(k+1)]_j$, for all $j = 1, \dots, p$. To adapt the proof, it suffice to note that the j -th component of $p \max_{i=1, \dots, m} \{ A_i x_i(\bar{\lambda}^*) - A_i [\bar{x}_{\text{LP}}^*]^{(i)} \}$ in (14) can be substituted with the sum of the p -largest values in the set $\{[A_i]_j x_i(\bar{\lambda}^*) - [A_i]_j [\bar{x}_{\text{LP}}^*]^{(i)}\}_{i=1}^m$, and the following derivations will remain unchanged. In the same vein one can redefine $\bar{\gamma}$ in (5) and change (18) in the proof of Theorem 4 below replacing $p \max_{i=1, \dots, m} \{ c_i^\top x_i(\bar{\lambda}^*) - c_i^\top [\bar{x}_{\text{LP}}^*]^{(i)} \}$ with the sum of the p -largest values in the set $\{c_i^\top x_i(\bar{\lambda}^*) - c_i^\top [\bar{x}_{\text{LP}}^*]^{(i)}\}_{i=1}^m$ to obtain a tighter bound also on the performance guarantees.

Proof of Theorem 4. Denote as $J_{\mathcal{P}}^*$, $J_{\bar{\mathcal{P}}_{\text{LP}}}^*$, and $J_{\mathcal{P}_{\text{LP}}}^*$ the optimal cost of \mathcal{P} , $\bar{\mathcal{P}}_{\text{LP}}$, and \mathcal{P}_{LP} , respectively. From Assumption 3 it follows that $J_{\mathcal{P}}^*$, $J_{\bar{\mathcal{P}}_{\text{LP}}}^*$, and $J_{\mathcal{P}_{\text{LP}}}^*$ are finite.

Consider the quantity $\sum_{i=1}^m c_i^\top x_i(k) - J_{\mathcal{P}}^*$. As in the proof of Theorem 3.3 in [18], we add and subtract $J_{\bar{\mathcal{P}}_{\text{LP}}}^*$ and $J_{\mathcal{P}_{\text{LP}}}^*$ to obtain

$$\begin{aligned}
\sum_{i=1}^m c_i^\top x_i(k) - J_{\mathcal{P}}^* &= \left(\sum_{i=1}^m c_i^\top x_i(k) - J_{\bar{\mathcal{P}}_{\text{LP}}}^* \right) \\
&\quad + (J_{\bar{\mathcal{P}}_{\text{LP}}}^* - J_{\mathcal{P}_{\text{LP}}}^*) + (J_{\mathcal{P}_{\text{LP}}}^* - J_{\mathcal{P}}^*). \quad (17)
\end{aligned}$$

We shall next derive a bound for each term in (17).

Bound on $\sum_{i=1}^m c_i^\top x_i(k) - J_{\mathcal{P}_{\text{LP}}}^$:*

Similarly to the proof of Theorem 3 for feasibility, due to Theorem 2.5 in [18], we have that there exists a set I of cardinality at least $m - p$ such that $x_i(\bar{\lambda}^*) = [\bar{x}_{\text{LP}}^*]^{(i)}$, for all $i \in I$. Therefore,

$$\begin{aligned} \sum_{i=1}^m c_i^\top x_i(\bar{\lambda}^*) - J_{\mathcal{P}_{\text{LP}}}^* &= \sum_{i=1}^m c_i^\top x_i(\bar{\lambda}^*) - \sum_{i=1}^m c_i^\top [\bar{x}_{\text{LP}}^*]^{(i)} \\ &= \sum_{i \in I^c} c_i^\top x_i(\bar{\lambda}^*) - c_i^\top [\bar{x}_{\text{LP}}^*]^{(i)} \\ &\leq p \max_{i=1, \dots, m} \left\{ c_i^\top x_i(\bar{\lambda}^*) - c_i^\top [\bar{x}_{\text{LP}}^*]^{(i)} \right\}, \quad (18) \end{aligned}$$

where $I^c = \{1, \dots, m\} \setminus I$.

According to [16, pag. 117], the components $[\bar{x}_{\text{LP}}^*]^{(i)}$ of the (unique, under Assumption 4) solution \bar{x}_{LP}^* to $\bar{\mathcal{P}}_{\text{LP}}$ is the limit point of the sequence $\{\tilde{x}_i(k)\}_{k \geq 1}$, defined as

$$\tilde{x}_i(k) = \frac{\sum_{r=1}^{k-1} \alpha(r) x_i(r+1)}{\sum_{r=1}^{k-1} \alpha(r)}.$$

By linearity, for all $k \geq 1$, we have that

$$c_i^\top \tilde{x}_i(k) = \frac{\sum_{r=1}^{k-1} \alpha(r) c_i^\top x_i(r+1)}{\sum_{r=1}^{k-1} \alpha(r)} \geq \min_{r \leq k} c_i^\top x_i(r) \geq \gamma_i,$$

where the first inequality is due to the fact that all $\alpha(k)$ are positive and the last one derives from the fact $\{\min_{r \leq k} c_i^\top x_i(r)\}_{k \geq 1}$ is a non-increasing sequence that takes values in a finite set, and hence is lower bounded by its limiting value γ_i . Therefore, by taking the limit for $k \rightarrow \infty$, we also have that

$$c_i^\top [\bar{x}_{\text{LP}}^*]^{(i)} \geq \gamma_i. \quad (19)$$

Since (18) holds for any choice of $x_i(\bar{\lambda}^*)$ which minimize $(c_i^\top + \bar{\lambda}^{*\top} A_i) x_i$ over $\text{vert}(X_i)$, by Proposition 6 it follows that, for $k \geq \bar{K}$, $x_i(\bar{\lambda}^*) = x_i(k)$ and, as a result

$$\begin{aligned} \sum_{i=1}^m c_i^\top x_i(k) - J_{\mathcal{P}_{\text{LP}}}^* &\leq p \max_{i=1, \dots, m} \left\{ c_i^\top x_i(k) - c_i^\top [\bar{x}_{\text{LP}}^*]^{(i)} \right\} \\ &\leq p \max_{i=1, \dots, m} \left\{ \max_{r \leq k} c_i^\top x_i(r) - c_i^\top [\bar{x}_{\text{LP}}^*]^{(i)} \right\} \\ &\leq p \max_{i=1, \dots, m} \left\{ \max_{r \leq k} c_i^\top x_i(r) - \gamma_i \right\}, \end{aligned}$$

where the second inequality is obtained by taking the maximum up to iteration k and the third inequality is due to (19).

Now if we recall the definition of $\gamma(k)$ in (5) and its finite-time convergence to $\bar{\gamma}$, jointly with the fact that γ_i is the limiting value of $\{\min_{r \leq k} c_i^\top x_i(r)\}_{k \geq 1}$, we finally get that there exists $K \geq \bar{K}$, such that for $k \geq K$

$$p \max_{i=1, \dots, m} \left\{ \max_{r \leq k} c_i^\top x_i(r) - \gamma_i \right\} = \bar{\gamma},$$

thus leading to

$$\sum_{i=1}^m c_i^\top x_i(k) - J_{\mathcal{P}_{\text{LP}}}^* \leq \bar{\gamma}, \quad k \geq K.$$

Bound on $J_{\bar{\mathcal{P}}_{\text{LP}}}^ - J_{\mathcal{P}_{\text{LP}}}^*$:*

Problem \mathcal{P}_{LP} can be considered as a perturbed version of $\bar{\mathcal{P}}_{\text{LP}}$, since the coupling constraint of $\bar{\mathcal{P}}_{\text{LP}}$ is given by

$$\sum_{i=1}^m A_i x_i \leq b - \bar{\rho}$$

and that of \mathcal{P}_{LP} can be obtained by adding $\bar{\rho}$ to its right-hand-side. From perturbation theory (see [7, Section 5.6.2]) it then follows that the optimal cost $J_{\mathcal{P}_{\text{LP}}}^*$ is related to $J_{\bar{\mathcal{P}}_{\text{LP}}}^*$ by:

$$J_{\bar{\mathcal{P}}_{\text{LP}}}^* - J_{\mathcal{P}_{\text{LP}}}^* \leq \bar{\lambda}^{*\top} \bar{\rho}. \quad (20)$$

From Assumption 5, by applying Lemma 1 in [14] we have that for all $\lambda \geq 0$

$$\|\bar{\lambda}^*\|_1 \leq \frac{1}{m\zeta} \left(\sum_{i=1}^m c_i^\top \hat{x}_i + \lambda^\top b - \sum_{i=1}^m \min_{x_i \in X_i} (c_i^\top + \lambda^\top A_i) x_i \right). \quad (21)$$

Substituting $\lambda = 0$ in (21) we get

$$\begin{aligned} \|\bar{\lambda}^*\|_1 &\leq \frac{1}{m\zeta} \left(\sum_{i=1}^m c_i^\top \hat{x}_i - \sum_{i=1}^m \min_{x_i \in X_i} c_i^\top x_i \right) \\ &\leq \frac{1}{\zeta} \max_{i=1, \dots, m} \left\{ \max_{x_i \in X_i} c_i^\top x_i - \min_{x_i \in X_i} c_i^\top x_i \right\} \\ &= \frac{\bar{\gamma}}{p\zeta}, \end{aligned} \quad (22)$$

where the second inequality comes from the fact that $c_i^\top \hat{x}_i \leq \max_{x_i \in X_i} c_i^\top x_i$ and that $\sum_{i=1}^m \beta_i \leq m \max_i \beta_i$ for any β_i , and the last equality is due to (3). Using (22)

in (20) we have

$$\begin{aligned} J_{\mathcal{P}_{LP}}^* - J_{\mathcal{P}_{LP}}^* &\leq \bar{\lambda}^{\star\top} \bar{\rho} \\ &\leq \|\bar{\lambda}^*\|_1 \|\bar{\rho}\|_\infty \\ &\leq \frac{\|\bar{\rho}\|_\infty}{p\zeta} \tilde{\gamma}, \end{aligned}$$

where the second inequality is due to the Hölder's inequality.

Bound on $J_{\mathcal{P}_{LP}}^ - J_{\mathcal{P}}^*$:*

Since \mathcal{P}_{LP} is a relaxed version of \mathcal{P} , then $J_{\mathcal{P}_{LP}}^* - J_{\mathcal{P}}^* \leq 0$.

The proof is concluded considering (17) and inserting the bounds obtained for the three terms. \square

4 Performance-oriented variant of Algorithm 1

While Algorithm 1 is able to find a feasible solution to \mathcal{P} , it does not directly consider the performance of the solution, whereas the user is concerned with both feasibility and performance with higher priority given to feasibility. This calls for a modification to Algorithm 1 which also takes into account the performance achieved.

Theorem 3 guarantees that there exists an iteration index K after which the iterates stay feasible for \mathcal{P} for all $k \geq K$. Now, suppose that the agents, together with the $A_i x_i(k)$ also transmit $c_i^\top x_i(k)$ to the central unit, then the central unit can construct the cost of $x(k) = [x_1(k)^\top, \dots, x_m(k)^\top]^\top$ at each iteration. When a feasible solution is found, its cost may be compared with that of a previously stored solution, and the central unit can decide to keep the new tentative solution or discard it. This way we are able to track the best feasible solution across iterations.

The modified procedure is summarized in Algorithm 2. Note that, compared to Algorithm 1, each agent is required to transmit also the cost of its tentative solution.

5 Application to optimal PEVs charging

In this section we show the efficacy of the proposed approach in comparison to the one described in [18] on the Plug-in Electric Vehicles (PEVs) charging problem described in [18]. This problem consists in finding an optimal overnight charging schedule for a fleet of m vehicles, which has to satisfy both local requirements and limitations (e.g., maximum charging power and desired final state of charge for each vehicle), and some network-wide constraints (i.e., maximum power that the network can deliver at each time slot). We consider both version of the PEVs charging problem, namely, the “charge only” setup in which all vehicles can only draw energy from

Algorithm 2 Performance-oriented version

```

1: % Initialize variables
2:  $\lambda \leftarrow 0, \bar{s}_i \leftarrow -\infty, \underline{s}_i \leftarrow +\infty, i = 1, \dots, m$ 
3:  $\tilde{J} \leftarrow +\infty, \delta \leftarrow 0, k \leftarrow 0$ 
4: repeat
5:   for  $i = 1$  to  $m$  do
6:     % Store tentative local solution
7:     if  $\delta = 1$  then
8:        $\tilde{x}_i \leftarrow x_i$ 
9:     end if
10:    % Update tentative local solution
11:     $x_i \leftarrow \arg \min_{x_i \in \text{vert}(X_i)} (c_i^\top + \lambda^\top A_i) x_i$ 
12:  end for
13:  % If solution is feasible and has better cost, then
  tell agents to update their tentative solutions
14:  if  $\sum_{i=1}^m A_i x_i \leq b$  and  $\sum_{i=1}^m c_i^\top x_i < \tilde{J}$  then
15:     $\tilde{J} \leftarrow \sum_{i=1}^m c_i^\top x_i$ 
16:     $\delta \leftarrow 1$ 
17:  else
18:     $\delta \leftarrow 0$ 
19:  end if
20:  % Update tightening
21:   $\bar{s}_i \leftarrow \max\{\bar{s}_i, A_i x_i\}, i = 1, \dots, m$ 
22:   $\underline{s}_i \leftarrow \min\{\underline{s}_i, A_i x_i\}, i = 1, \dots, m$ 
23:   $\rho_i \leftarrow \bar{s}_i - \underline{s}_i, i = 1, \dots, m$ 
24:   $\rho \leftarrow p \max\{\rho_1, \dots, \rho_m\}$ 
25:  % Update dual variables
26:   $\lambda \leftarrow [\lambda + \alpha(k)(\sum_{i=1}^m A_i x_i - b + \rho)]_+$ 
27:  % Update iteration counter
28:   $k \leftarrow k + 1$ 
29: until time is over

```

the network, and the “vehicle to grid” setup where the vehicles are also allowed to inject energy in the network.

The improvement of our approach with respect to that in [18] is measured in terms of the following two relative indices: the reduction in the level of conservativeness $\Delta\rho\%$ and the improvement in performance achieved by the primal solution $\Delta J\%$ defined as

$$\Delta\rho\% = \frac{\|\tilde{\rho}\|_\infty - \|\bar{\rho}\|_\infty}{\|\tilde{\rho}\|_\infty} \cdot 100 \quad \text{and} \quad \Delta J\% = \frac{J_{\tilde{\rho}} - J_{\bar{\rho}}}{J_{\tilde{\rho}}} \cdot 100,$$

where $J_{\tilde{\rho}} = \sum_{i=1}^m c_i^\top x_i(\lambda_{\tilde{\rho}}^*)$ and $J_{\bar{\rho}} = \sum_{i=1}^m c_i^\top x_i(\bar{\lambda}^*)$. A positive value for these indices indicates that our approach is less conservative.

For a thorough comparison we determined the two indices while varying: i) the number of vehicles in the network, ii) the realizations of the random parameters entering the system description (cost of the electrical energy and local constraints), and iii) the right hand side of the joint constraints. All parameters and their probability distributions were taken from Table 1 in [18].

In Table 1 we report the conservativeness reduction and the cost improvement for the “vehicle to grid” setup. As




Fig. 1. Histogram of the performance improvement ($\Delta J\%$) achieved by the primal solution obtained by the proposed method with respect to the one proposed in [18] over 1000 runs.

it can be seen from the table, the level of conservativeness is reduced by 50% while the improvement in performance (witnessed by positive values of $\Delta J\%$) drops as the number of agents grows. This is due to the fact that the relative gap between $J_{\tilde{\rho}}$ and $J_{\mathcal{P}}^*$ tends to zero as $m \rightarrow \infty$, thus reducing the relative margin for performance improvement.

We do not report the results for the “charge only” setup since the two methods lead to the same level of conservativeness and performance of the primal solution.

We also tested the proposed approach against changes of the random parameters defining the problem. We fixed $m = 250$ and performed 1000 tests running Algorithm 1 and the approach in [18] with different realization for all parameters, extracted independently. Figure 1 plots an histogram of the values obtained for $\Delta J\%$ in the 1000 tests. Note that the cost improvement ranges from 3% to 15% and, accordingly to the theory, is always non-negative. The reduction in the level of conservativeness is also in this case 50%, suggesting that the proposed iterative scheme exploits some structure in the PEVs charging problem that the approach in [18] overlooks. Also in this case, in the “charge only” setup the two methods lead to the same level of conservativeness and performance.

Finally, we compared the two approaches in the “vehicle to grid” setup against changes in the joint constraints. If the number of electric vehicles is $m = 250$ and we decrease the maximum power that the network can deliver

m	250	500	1000	2500	5000	10000
$\Delta \rho\%$	50%	50%	50%	50%	50%	50%
$\Delta J\%$	13.9%	3.1%	1.1%	0.15%	0.05%	0.02%

Table 1

Reduction in the level of conservativeness ($\Delta \rho\%$) and improvement in performance ($\Delta J\%$) achieved by the primal solution obtained by the proposed method when compared with the one proposed in [18].




Fig. 2. Histogram of the relative distance from the optimal value of \mathcal{P} achieved by the primal solution obtained by Algorithm 1 (blue) and Algorithm 2 (orange), over 1000 runs.

by 37%, then the $\tilde{\rho}$ that results from applying the approach in [18] makes $\tilde{\mathcal{P}}_{LP}$ infeasible, thus violating Assumption 1. Whereas $\bar{\mathcal{P}}_{LP}$ associated with the limiting value $\bar{\rho}$ for $\{\rho(k)\}_{k \geq 1}$ in Algorithm 1 remains feasible.

5.1 Comparison between Algorithms 1 and 2

To show the benefits of Algorithm 2 in terms of performance, we run 1000 test with $m = 250$ vehicles in the “charge only” setup, where we are also able to compute the optimal solution of \mathcal{P} , and compare the performance of Algorithm 1 and 2 in terms of relative distance from the optimal cost $J_{\mathcal{P}}^*$ of \mathcal{P} .

Figure 2 shows the distribution of $(J_{\tilde{\rho}} - J_{\mathcal{P}}^*)/J_{\mathcal{P}}^* \cdot 100$ obtained with Algorithm 1 (blue) and $(\tilde{J} - J_{\mathcal{P}}^*)/J_{\mathcal{P}}^* \cdot 100$ obtained with Algorithm 2 (orange) for the 1000 runs. As can be seen from the picture, most runs of Algorithm 2 result in a performance very close to the optimal one, while the runs from Algorithm 1 exhibit lower performance.

6 Concluding remarks

We proposed a new method for computing a feasible solution to a large-scale mixed integer linear program via a decentralized iterative scheme that decomposes the program in smaller ones and has the additional beneficial side-effect of preserving privacy of the local information if the problem originates from a multi-agent system.

This work improves over existing state-of-the-art results in that feasibility is achieved in a finite number of iterations and the decentralized solution is accompanied by a less conservative performance guarantee. The application to a plug-in electric vehicles optimal charging problem verifies the improvement gained in terms of performance.

Our method was recently extended to a distributed setup without any central unit in [10], by integrating within

the decentralized iterative scheme proposed here a max-consensus algorithm on the tightening coefficient and employing the distributed approach for updating the dual variables proposed in [9]. Finite convergence properties are retained in the distributed scheme.

Future research directions include the computation of an upper bound on the number of iterations needed for convergence. This is more critical in a distributed setup where no central unit exists that can directly monitor feasibility and/or inspect performance. Moreover, we aim at exploiting the analysis of [17] to generalize our results to problems with nonconvex objective functions.

A Appendix

We provide an example illustrating that the two solution methodologies outlined in Section 1, namely, using the optimal dual solution to recover a primal one via (1), and using a subgradient methodology together with the averaging procedure described in [16, pag. 117], may both lead to infeasible solutions.

Example 1 Consider the following problem

$$\begin{aligned} \min_x \quad & x \\ \text{subject to:} \quad & -x \leq -0.5 \\ & x \in \{0, 1, 2\}, \end{aligned}$$

whose dual (dualizing only the inequality constraint) is given by

$$\max_{\lambda \geq 0} \quad 0.5\lambda + \min_{x \in \{0, 1, 2\}} (1 - \lambda)x.$$

We will now apply the two solution methodologies outlined in Section 1.

Part 1 If we solve the dual up to optimality we get $\lambda^* = 1$ as the unique solution. Using (1) with $\lambda = 1$ we get $x(\lambda^*) = 0$ and $x(\lambda^*) = 2$ as possible solutions. Clearly, $x(\lambda^*) = 0$ is feasible for the local constraint but it is not feasible for the dualized constraint $x \geq 0.5$.

Part 2 If we employ the averaging procedure described in [16, pag. 117] while solving the dual using the subgradient method, then we will converge to the unique solution of the following convexified program

$$\begin{aligned} \min_x \quad & x \\ \text{subject to:} \quad & -x \leq -0.5 \\ & x \in \text{conv}(\{0, 1, 2\}) = [0, 2], \end{aligned}$$

which is $x_{\text{LP}}^* = 0.5$. Clearly, x_{LP}^* satisfies the dualized constraint but not the local one.

References

- [1] J.-P. Aubin and I. Ekeland. Estimates of the duality gap in nonconvex optimization. *Mathematics of Operations Research*, 1(3):225–245, 1976.
- [2] P. Baumann and N. Trautmann. Portfolio-optimization models for small investors. *Mathematical Methods of Operations Research*, 77(3):345–356, 2013.
- [3] A. Bemporad and M. Morari. Control of systems integrating logic, dynamics, and constraints. *Automatica*, 35(3):407–427, 1999.
- [4] D. Bertsekas, G. Lauer, N. Sandell, and T. Posbergh. Optimal short-term scheduling of large-scale power systems. *IEEE Transactions on Automatic Control*, 28(1):1–11, 1983.
- [5] D. P. Bertsekas. *Nonlinear programming*. Athena scientific Belmont, 1999.
- [6] D. Bertsimas and J. N. Tsitsiklis. *Introduction to linear optimization*, volume 6. Athena Scientific Belmont, MA, 1997.
- [7] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [8] M. Dawande, S. Gavirneni, and S. Tayur. Effective heuristics for multiproduct partial shipment models. *Operations research*, 54(2):337–352, 2006.
- [9] A. Falsone, K. Margellos, S. Garatti, and M. Prandini. Dual decomposition for multi-agent distributed optimization with coupling constraints. *Automatica*, 84:149–158, 2017.
- [10] A. Falsone, K. Margellos, and M. Prandini. A distributed iterative algorithm for multi-agent MILPs: Finite-time feasibility and performance characterization. *IEEE Control Systems Letters*, 2(4):563–568, 2018.
- [11] A. M. Geoffrion. Lagrangean relaxation for integer programming. *Mathematical Programming Study* 2, pages 82–114, 1974.
- [12] D. Ioli, A. Falsone, and M. Prandini. An iterative scheme to hierarchically structured optimal energy management of a microgrid. In *54th Conference on Decision and Control (CDC2015)*, pages 5227–5232, 2015.
- [13] M. Jünger, T. M. Liebling, D. Naddef, G. L. Nemhauser, W. R. Pulleyblank, G. Reinelt, G. Rinaldi, and L. A. Wolsey. *50 Years of Integer Programming 1958-2008: From the Early Years to the State-of-the-art*. Springer Science & Business Media, 2009.
- [14] A. Nedić and A. Ozdaglar. Approximate primal solutions and rate analysis for dual subgradient methods. *SIAM Journal on Optimization*, 19(4):1757–1780, 2009.
- [15] N. J. Redondo and A. Conejo. Short-term hydro-thermal coordination by lagrangian relaxation: solution of the dual problem. *IEEE Transactions on Power Systems*, 14(1):89–95, 1999.
- [16] N. Shor. *Minimization Methods for Non-Differentiable Functions*. Springer, 1985.
- [17] M. Udell and S. Boyd. Bounding duality gap for separable problems with linear constraints. *Computational Optimization and Applications*, 64(2):355–378, 2016.
- [18] R. Vujanic, P. M. Esfahani, P. J. Goulart, S. Mariéthoz, and M. Morari. A decomposition method for large scale MILPs, with performance guarantees and a power system application. *Automatica*, 67:144–156, 2016.
- [19] H. Y. Yamin. Review on methods of generation scheduling in electric power systems. *Electric Power Systems Research*, 69(2):227–248, 2004.