



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE DELLA RICERCA

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

Distributed Mixed-Integer Linear Programming via Cut Generation and Constraint Exchange

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Testa A., Rucco A., Notarstefano G. (2020). Distributed Mixed-Integer Linear Programming via Cut Generation and Constraint Exchange. IEEE TRANSACTIONS ON AUTOMATIC CONTROL, 65(4), 1456-1467 [10.1109/TAC.2019.2920812].

Availability:

This version is available at: <https://hdl.handle.net/11585/790159> since: 2021-01-21

Published:

DOI: <http://doi.org/10.1109/TAC.2019.2920812>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the post peer-review accepted manuscript of:

A. Testa, A. Rucco and G. Notarstefano, "Distributed Mixed-Integer Linear Programming via Cut Generation and Constraint Exchange," in IEEE Transactions on Automatic Control, vol. 65, no. 4, pp. 1456-1467, April 2020.

The published version is available online at:

<https://doi.org/10.1109/TAC.2019.2920812>

© 2016 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Distributed Mixed-Integer Linear Programming via Cut Generation and Constraint Exchange

Andrea Testa¹, Alessandro Rucco¹, and Giuseppe Notarstefano²

Abstract—Many problems of interest for cyber-physical network systems can be formulated as Mixed-Integer Linear Programs in which the constraints are distributed among the agents. In this paper we propose a distributed algorithmic framework to solve this class of optimization problems in a peer-to-peer network with no coordinator and with limited computation and communication capabilities. At each communication round, agents locally solve a small linear program, generate suitable cutting planes and communicate a fixed number of active constraints. Within the distributed framework, we first propose an algorithm that, under the assumption of integer-valued optimal cost, guarantees finite-time convergence to an optimal solution. Second, we propose an algorithm for general problems that provides a suboptimal solution up to a given tolerance in a finite number of communication rounds. Both algorithms work under asynchronous, directed, unreliable networks. Finally, through numerical computations, we analyze the algorithm scalability in terms of the network size. Moreover, for a multi-agent multi-task assignment problem, we show, consistently with the theory, its robustness to packet loss.

I. INTRODUCTION

Mixed-Integer Linear Programs (MILPs) play an important role in several control problems, including control of hybrid systems [2], trajectory planning [3], and task assignment [4]. For example, since non-convex functions can be approximated by means of piecewise-linear functions, see, e.g., [2], nonlinear optimal control problems can be approximated by MILPs. Though MILPs are known to be \mathcal{NP} -hard, numerous efficient algorithms exist in a centralized setup. A widely used approach is the *cutting-plane* method, see, e.g. [5], which is based on the iterative solution of linear programming relaxations (obtained by neglecting integer constraints) and the generation of linear constraints (cuts) removing the current non-integer solution from the feasible set. In this paper, we consider the following *Mixed-Integer Linear Program* (MILP)

$$\begin{aligned} & \min_z c^\top z \\ & \text{subj. to } a_i^\top z \leq b_i, i = 1, \dots, n \\ & \quad z \in \mathbb{Z}^{d_Z} \times \mathbb{R}^{d_R}, \end{aligned} \quad (1)$$

¹A. Testa and A. Rucco are with the Department of Engineering, Università del Salento, Lecce, Italy, name.lastname@unisalento.it.

²G. Notarstefano is with the Department of Electrical, Electronic and Information Engineering, University of Bologna, Bologna, Italy, giuseppe.notarstefano@unibo.it.

An early short version of this work appeared as [1]: the current article includes a much improved comprehensive treatment, an extended algorithm for general mixed-integer linear programs, new numerical computations, and an application to the multi-agent multi-task assignment problem.

This result is part of a project that has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 638992 - OPT4SMART).

where d_Z and d_R are the dimensions of the integer and real variables, $d = d_Z + d_R$, $a_i \in \mathbb{R}^d$, $b_i \in \mathbb{R}$, $c \in \mathbb{R}^d$, and n is the number of inequality constraints.

We deal with a distributed setup in which MILP constraints are shared among agents of a network, which aim at solving the whole problem by local computations and communications. An important challenge that needs to be taken into account in a distributed context is that the communication can be asynchronous, unreliable, and the topology directed. Inspired by the centralized literature, we will propose a distributed algorithmic framework based on cutting planes.

Only few works address the solution of MILPs in a purely distributed way, that is, with peer processors communicating over a network, without the presence of a central (master) unit. For this reason we organize the relevant literature to our paper in two main blocks: centralized and parallel approaches to solve MILPs in control applications, and distributed algorithms solving Linear Programs (LPs) or convex programs arising as relaxations or special versions of MILPs. As for centralized approaches, a Model Predictive Control (MPC) scheme to solve constrained multivariable control problems is proposed in [6], [7]. The MPC is formulated as a multi-parametric MILP for which solver [8] is used. In [9] a branch-and-bound procedure is devised for the computation of optimal and suboptimal solutions to parametric MILPs. In [3] and [10], a collision-free trajectory optimization problem for autonomous vehicles is formulated as a MILP solved by a branch-and-bound algorithm with branching heuristics. In [11], a multi-robot routing problem under connectivity constraints is shown to be formulated as an integer program with binary variables, and then its LP relaxation is solved. Recently, in [12] a heuristic based on the Alternating Direction Method of Multipliers is used to approximately solve mixed-integer linear and quadratic programs. The heuristic is applied to the control of hybrid vehicles. As for parallel methods, in [13] a Lagrange relaxation approach is used in order to decompose the overall MILP into multiple subproblems each of which is solved in a client-server parallel architecture. The proposed solution is applied to the demand response control in smart grids. In [14] a parallel dual decomposition method, relying on a suitable tightening of the constraints, is proposed to approximately solve structured MILPs with local and coupling constraints. The algorithm is improved in [15] by means of an iterative tightening procedure. The methods are applied to charging control of electric vehicles.

As for distributed optimization algorithms, we concentrate our review on schemes solving linear programs or convex programs that represent a relaxation of common mixed-integer

programs. In [16] a robust, distributed algorithm is designed to solve linear programs over networks with event-triggered communication. In [17] a distributed algorithm is proposed to find valid solutions for the so called bargaining problem, which is an integer program, by means of a linear program relaxation. In [18] a (mixed-integer) utility maximization problem is addressed. The proposed solution is based on a convex relaxation obtained by neglecting the integer constraint on the rates. In [19] and [20] the authors propose a Newton-type fast converging algorithm to solve Network Utility Maximization problems with self-concordant utility functions. In [21] consensus algorithms are proposed to solve abstract optimization programs (i.e., a generalization of LPs) in asynchronous networks, while a distributed simplex algorithm is proposed in [22] to solve degenerate LPs and multi-agent assignment problems. A distributed version of the Hungarian method is proposed in [23] to solve LPs arising in multi-robot assignment problems. In [24] and [25] approximate solutions for task assignment (MILP) problems are proposed based respectively on a simplex ascent and an auction approach. To conclude, [26], [27], [28] are first attempts of proposing a distributed solution for MILPs. We discuss the main differences with our approach after the contributions paragraph.

The contributions of this paper are as follows. We propose a distributed algorithmic framework, named Distributed Cutting Plane and Constraint Exchange for MILPs (DiCUT-MILP), based on the local generation of cutting planes, the solution of local LP relaxations of (1), and the exchange of active constraints. Specifically, we propose two distributed algorithms. The first one, called INT-DiCUT-MILP, has guaranteed finite-time convergence to an optimal solution of (1) under the assumption of integer-valued optimal cost. All relevant Integer Programs can, e.g., be casted in this setup. To remove the assumption of integer-valued optimal cost, we then propose an algorithm, called ϵ -DiCUT-MILP, “practically” solving (1), i.e., computing a feasible point with cost exceeding the optimal one no more than ϵ . Both algorithms involve, as local computations at each node, only a LP solver (processing a number of constraints depending only on the problem dimension and the number of neighboring nodes) and a simple procedure for the generation of cutting planes. Thus, the algorithms are scalable in terms of local memory, computation, and communication. Moreover, under slightly stronger assumptions on the graph, we provide a halting condition allowing agents to stop the algorithm in a purely distributed way. To the best of our knowledge, these are the first distributed algorithms solving MILPs in a distributed context. Notably, the proposed algorithms work under asynchronous, unreliable, and directed networks, so that they are immediately implementable in concrete scenarios.

We highlight some meaningful differences with respect to the literature discussed above. In [14] and [15], suboptimal algorithms with performance guarantees for MILPs are given and a central unit is required. Papers [26] and [27] propose strategies to find feasible (suboptimal) points for problem (1). Moreover, in [27], agents perform the local computation in a sequential order, while in [26] a gossip protocol is considered in which one node per time becomes active. In [28] a distributed algorithm with performance guarantees, based on

dual decomposition and a time-varying restriction technique, is proposed. The algorithm applies to a class of optimization problems in which agents aim at minimizing the sum of local linear cost functions, subject to local linear constraints, while all local variables are coupled by global constraints. Our algorithm is developed for a different class of MILPs with common cost and local constraints. The algorithm in [28] converges to a suboptimal solution with a performance guarantee that depends on the problem data. Our approach, instead, finds a suboptimal solution with an arbitrary tolerance and works over a general, possibly unreliable, asynchronous network. Finally, although related approaches have been proposed in [21] and [29] for problems with continuous decision variables, novel tools are needed in this paper due to the mixed-integer nature of the problem.

The paper is organized as follows. In Section II we recall the centralized cutting-plane approach for MILPs. In Section III we introduce a distributed meta-algorithm together with two specific algorithms, while in Section IV we analyze their convergence. Numerical computations are provided in Section V for randomly generated MILPs and for a multi-agent multi-task assignment setup.

Notation: Given the decision variable $z \in \mathbb{Z}^{dz} \times \mathbb{R}^{dr}$ of (1), we denote by $x \in \mathbb{Z}^{dz}$ the vector of variables subject to integer constraints, and by $y \in \mathbb{R}^{dr}$ the vector of variables not required to be integer. We denote by e_ℓ the ℓ -th vector of the canonical basis (e.g., $e_1 = [1 \ 0 \ \dots \ 0]^\top$) of proper dimension. Given a vector $v \in \mathbb{R}^d$, we denote by v_ℓ the ℓ -th component of v . Given two vectors $v, w \in \mathbb{R}^d$, v is lexicographically greater than w , $v >_{lex} w$, if there exists $\ell \in \{1, \dots, d\}$ such that $v_\ell > w_\ell$ and $v_m = w_m$ for all $m < \ell$. Given an inequality $a^\top z \leq b$ for $z \in \mathbb{R}^d$, with $a \in \mathbb{R}^d$ and $b \in \mathbb{R}$, we use the following simplified notation $\{a^\top z \leq b\} := \{z \in \mathbb{R}^d : a^\top z \leq b\}$ for the related half-space. The polyhedron induced by the inequality constraints $a_i^\top z \leq b_i$, $i = 1, \dots, n$, is $P := \bigcap_{i=1}^n \{a_i^\top z \leq b_i\}$. Recall that a polyhedron is a set described by the intersection of a finite number of half-spaces. In this paper we assume a LP solver is available. In particular, we use the simplex algorithm proposed in [30] to find the unique *lexicographically minimal optimal* (*lex-optimal* for short) solution of degenerate linear programs. From now on, we call such a solver LPLEXSOLV and say that it returns the lex-optimal solution of the solved LP. LPLEXSOLV also returns an optimal *basis* identifying the lex-optimal solution. Given a LP with constraint set $P := \bigcap_{i=1}^n P_i$, with each P_i a half-space, a basis B is the intersection of a *minimal* number of half-spaces $P_{\ell_1}, \dots, P_{\ell_q}$, $q \leq d$, such that the solution of the LP over the constraint set B is the same as the one over P . If the lex-optimal solution is considered, it turns out that B is the intersection of exactly d half-spaces.

II. CUTTING-PLANE FRAMEWORK FOR MIXED-INTEGER LINEAR PROGRAMMING

In this section we provide a brief description of one of the most used centralized methods to solve a MILP, i.e., the cutting-plane approach, see e.g., [31]. We introduce suitable cutting planes, namely intersection cuts and cost-based cuts, and describe centralized algorithms for MILPs, based on these cuts, that are relevant for our distributed framework.

A. Cutting-Plane approach for MILP

Let $P_I := P \cap (\mathbb{Z}^{d_z} \times \mathbb{R}^{d_R})$ be the set of feasible points of MILP (1), also called *mixed-integer set*. The optimal solutions of (1) are also optimal solutions of the following LP, [32],

$$\begin{aligned} & \min_z c^\top z \\ & \text{subj. to } z \in \text{conv}(P_I) \end{aligned} \quad (2)$$

where $\text{conv}(P_I)$ is the *convex hull* of P_I . A two-dimensional representation of a mixed-integer set and its convex hull is given in Figure 1.




Fig. 1. Example of a polyhedron P (gray area), with feasible set P_I (solid red lines) with $x \in \mathbb{Z}$, $y \in \mathbb{R}$ and its convex hull (blue striped area).

It is worth noting that, if P is a bounded polyhedron, by Meyer's Theorem, [33], $\text{conv}(P_I)$ is a polyhedron. For this reason, we make the following assumption, which is common in MILP literature.

Assumption 2.1 (Boundedness and Feasibility): The polyhedron P is bounded and $\text{conv}(P_I)$ is nonempty. \square

The main idea of the cutting-plane approach for MILPs is to neglect the integer constraints on the decision vector x (i.e., $x \in \mathbb{R}^{d_z}$), and iteratively solve *relaxed* linear problems in which the polyhedron P is tightened by additional half-spaces called *cutting planes* or *cuts*. The procedure terminates when the solution of the LP relaxation, call it $z^{\text{LP}} = (x^{\text{LP}}, y^{\text{LP}})$, is in $\text{conv}(P_I)$ with $x^{\text{LP}} \in \mathbb{Z}^{d_z}$. A valid cutting plane is a half-space containing $\text{conv}(P_I)$ but not z^{LP} . We introduce a CUTORACLE subroutine defined as follows. CUTORACLE(z^{LP}, P, c) returns the intersection of p cutting planes, $\{\alpha^\top z \leq \beta\}$ with $\alpha \in \mathbb{R}^{p \times d}$ and $\beta \in \mathbb{R}^p$, or \mathbb{R}^d . An iterative cutting-plane scheme can be recast in the form of the following meta-algorithm.

Centralized Cutting-Plane Meta-Algorithm:

1. *Initialization:* set $P = \bigcap_{i=1}^n \{a_i z \leq b_i\}$.
2. *LP solver:* Find an optimal solution $z^{\text{LP}} = (x^{\text{LP}}, y^{\text{LP}})$ of the LP relaxation of (1) with polyhedron P .
3. *Check feasibility:* if $x^{\text{LP}} \in \mathbb{Z}^{d_z}$, go to 6.
4. *Cutting-Plane:* $h = \text{CUTORACLE}(z^{\text{LP}}, P, c)$.
5. *Update:* $P = P \cap h$ and go to 2.
6. *Output:* z^{LP} .

It is worth noting that, at each iteration, the meta-algorithm uses the entire set of inequality constraints, P , and all the cuts generated up to that iteration.

Following the meta-algorithm, numerous algorithms have been proposed in the literature that have different convergence properties depending on CUTORACLE and the problem structure. In the next subsections we describe two centralized

algorithms. The key ingredients of both algorithms are: i) *Mixed-Integer Gomory* (MIG) cuts, [34], and ii) *cost-based cuts* generated according to the current cost value $c^\top z^{\text{LP}}$.

B. Gomory's Cutting-Plane Algorithm

Mixed-Integer Gomory (MIG) cuts are cutting planes proposed in [34] for MILPs with integer-valued optimal cost. Next, we introduce the notions of split disjunction, [35], and intersection cut, [36], that are intimately related to MIG cuts.

Definition 2.2 (Split Disjunction [35]): Given $\pi \in \mathbb{Z}^{d_z}$ and $\pi_0 \in \mathbb{Z}$, a split disjunction $D(\pi, \pi_0)$ is a set of the form $D(\pi, \pi_0) := \{\pi^\top x \leq \pi_0\} \cup \{\pi^\top x \geq \pi_0 + 1\}$. \square

Let B^{LP} be a basis for the optimal solution $z^{\text{LP}} = (x^{\text{LP}}, y^{\text{LP}})$ of a given LP relaxation of (1), and $D(\pi, \pi_0)$ a disjunction with respect to x^{LP} . Let $C(z^{\text{LP}})$ be the translated simplicial cone¹ formed by the intersection of the half-planes defining B^{LP} with apex in z^{LP} . *Intersection cuts* can be derived by considering the intersection between the extreme rays of $C(z^{\text{LP}})$ and the hyperplanes defining the split disjunction $D(\pi, \pi_0)$. A more detailed definition can be found in [36]. A two dimensional representation of a split disjunction with $\pi = [1 \ 0]^T$ is given in Figure 2 together with the corresponding intersection cut (dashed blue line) for a given basis of z^{LP} (solid lines).




Fig. 2. Example of split disjunction (gray shaded area) and intersection cut (dashed blue line) with respect to the basis B^{LP} (solid lines) in \mathbb{R}^2 .

It can be shown that the MIG cut with respect to x_ℓ^{LP} , i.e., the ℓ -th component of x^{LP} , is the intersection cut to the split disjunction $D(e_\ell, \lfloor x_\ell^{\text{LP}} \rfloor)$ and the basis B^{LP} , and is a valid cutting plane [37]. Next we show how MIG cuts can be generated. Let $z^{\text{LP}} = (x^{\text{LP}}, y^{\text{LP}})$ be the lex-optimal solution of a generic LP relaxation of problem (1) with $x_\ell^{\text{LP}} \notin \mathbb{Z}$ for some $\ell \in \{1, \dots, d_z\}$, and let $B^{\text{LP}} = \{A_B z \leq b_B\}$ be an associated basis. Consider the split disjunction $D(e_\ell, \lfloor x_\ell^{\text{LP}} \rfloor)$ which satisfies $z^{\text{LP}} \notin D(e_\ell, \lfloor x_\ell^{\text{LP}} \rfloor)$. Let r^m be the m -th column of $-A_B^{-1}$ and define λ_m as $\frac{\lfloor z_\ell^{\text{LP}} \rfloor - e_\ell^\top z^{\text{LP}}}{e_\ell^\top r^m}$ if $e_\ell^\top r^m < 0$, $\frac{\lfloor z_\ell^{\text{LP}} \rfloor - e_\ell^\top z^{\text{LP}} + 1}{e_\ell^\top r^m}$ if $e_\ell^\top r^m > 0$ and ∞ if $e_\ell^\top r^m = 0$. As it will be explained in Section IV, r^m represents an extreme ray (originating at z^{LP}) intersecting the disjunction, while λ_m is the displacement along the extreme ray of the intersection point. Let $\Lambda \in \mathbb{R}^d$ be the vector with m -th entry $\Lambda_m = 1/\lambda_m$ ($\Lambda_m = 0$ if $\lambda_m = \infty$). Then, the intersection cut to the split disjunction $D(e_\ell, \lfloor x_\ell^{\text{LP}} \rfloor)$ and the basis B^{LP} is $h_{\text{MIG}} = \{(\Lambda^\top A_B)^\top z \leq \Lambda^\top b_B - 1\}$.

Let us consider now the first non-integer component of x^{LP} , namely $x_{k_{\text{lex}}}^{\text{LP}}$ where $k_{\text{lex}} = \min\{k = 1, \dots, d_z : x_k^{\text{LP}} \notin \mathbb{Z}\}$. We call MIGORACLE the oracle that generates the MIG cut h_{MIG} with respect to $x_{k_{\text{lex}}}^{\text{LP}}$. We highlight that MIGORACLE returns h_{MIG} only when a variable subject to integer constraint

¹Given a cone $S \subset \mathbb{R}^d$ and a point $p \in \mathbb{R}^d$, the set $p + S$ is a translated cone with apex in p .

is not integer, i.e., $x^{\text{LP}} \notin \mathbb{Z}^{d_z}$. If the optimal solution z^{LP} obtained by solving the LP is such that $x^{\text{LP}} \in \mathbb{Z}^{d_z}$, then MIGORACLE does not return any constraint.

It is worth noting that if one implements the Centralized Cutting-Plane Meta-Algorithm using MIGORACLE as CUTORACLE, the algorithm may not converge. Indeed, a “tailing-off” phenomenon can be observed: a large number of cuts may be added without significant improvement in the cost. A simple two-dimensional example is discussed in [38].

We are now ready to describe Gomory’s Cutting-Plane Algorithm for MILPs with integer-valued optimal cost [34]. We provide a reformulation, given in [37], for MILPs in the form (1). Given a basis B^{LP} identifying the optimal solution z^{LP} and the corresponding cost function value $c^\top z^{\text{LP}}$, we define CUTORACLE as follows

$$(h_{\text{MIG}}, h_c) = \text{CUTORACLE}(z^{\text{LP}}, B^{\text{LP}}, c)$$

where h_{MIG} is the MIG cut generated by MIGORACLE and $h_c = \{c^\top z \geq \lceil c^\top z^{\text{LP}} \rceil\}$. It can be shown that, if the optimal objective function value is integer, Gomory’s cutting plane algorithm converges to an optimal solution in a finite number of iterations [34].

C. Algorithms for ϵ -suboptimal solutions

Another relevant piece of literature regards algorithms solving MILPs up to an arbitrary tolerance, namely based on the following notion of suboptimal solution.

Definition 2.3 (ϵ -suboptimal solution): Given a MILP as in (1), we say that z^ϵ is an ϵ -suboptimal solution to (1) if z^ϵ is feasible and satisfies $c^\top z^\epsilon - c^\top z^* \leq \epsilon$, where z^* is an optimal solution to (1). \square

In [38] the authors propose a cutting-plane algorithm, based on “variable” disjunctions, converging to an ϵ -suboptimal solution of the MILP in a finite number of iterations. In contrast to the usual cutting-plane approaches, which generate constraints only at the optimal solution of the current LP relaxation, the proposed algorithm generates constraints at multiple, near-optimal vertices. In [37] the author proposes an approximation algorithm which is not a classical cutting-plane method because it generates cuts that might not be valid for the mixed-integer set. Moreover, the algorithm relies on an inner procedure to check if feasible points have been cut off, which involves the solution of a MILP with integer-valued optimal cost.

III. A DISTRIBUTED CUTTING-PLANE AND CONSTRAINT EXCHANGE APPROACH FOR MILPs

Inspired by the centralized cutting-plane meta-algorithm, in this section we first propose a distributed meta-algorithm, called DiCUT-MILP, based on the local generation of cutting planes and the exchange of active constraints. As in the centralized case, it is not reasonable to provide a general, unified convergence analysis for the meta-algorithm, but proper tools are needed for specific algorithms. Thus, based on the high-level methodological approach, we provide an algorithm for MILPs with integer-valued optimal cost (as in [34]) and an approximation-based algorithm for general problems. Then,

we discuss some key features of the proposed algorithms and provide a distributed stopping criterion. We first formalize the distributed computation setup.

A. Distributed Optimization Setup

In our distributed setup, we consider a network composed by a set of agents $V = \{1, \dots, N\}$. In general, the $n \geq N$ constraints in problem (1) are distributed among the agents, so that each agent knows only a small number of constraints. For simplicity, we assume one constraint $\{a_i^\top z \leq b_i\}$ is assigned to the i -th agent, so that $N = n$, but we will keep the two notations separated to show that the algorithm can be easily implemented also when $n > N$ (i.e., more than one constraint is assigned to agents). The communication among the agents is modeled by a time-varying digraph $\mathcal{G}_c(t) = (V, E(t))$, with $t \in \mathbb{N}$ being a universal slotted time. A digraph $\mathcal{G}_c(t)$ models the communication in the sense that there is an edge $(i, j) \in E(t)$ if and only if agent i is able to send information to agent j at time t . For each node i , the set of *in-neighbors* of i at time t is denoted by $\mathcal{N}_i(t)$ and is the set of all j such that there exists an edge $(j, i) \in E(t)$. A static digraph is *strongly connected* if there exist a directed path for each pair of agents i and j . For a static digraph $\mathcal{G}_c = (V, E)$, we use $d_{\mathcal{G}}$ to denote the graph diameter, that is the maximum distance taken over all the pairs of agents (i, j) , where the distance is defined as the length of the shortest directed path from i to j . A time-varying digraph is *jointly strongly connected* if, for all $t \in \mathbb{N}$, $\cup_{\tau=t}^{\infty} \mathcal{G}_c(\tau)$ is strongly connected. The time-varying digraph is said to be (*uniformly jointly*) *L-strongly connected* if there exists an integer $L \geq 1$ such that, for all $t \in \mathbb{N}$, the graph $\cup_{\tau=t}^{t+L-1} \mathcal{G}_c(\tau)$ is strongly connected.

B. Meta-Algorithm description

We now describe the proposed distributed meta-algorithm. In contrast to the centralized approach, in the distributed setup, some agents may need to solve local LP relaxations which are unbounded, especially at the first iterations. For this reason, we initialize the algorithm by assigning to each agent a set of artificial constraints which are inactive for problem (1). This method is often referred to as big-M method. Specifically, the decision variable of each agent is delimited by a box constraint. In particular, for a given, sufficiently large $M > 0$, we define the bounding box $H_M := \bigcap_{\ell=1}^d (\{z_\ell \leq M\} \cap \{z_\ell \geq -M\})$.

In the analysis we will need the following assumption.

Assumption 3.1 (Bounding Box): Given a MILP with polyhedron P and a bounding box H_M , then $P \subseteq H_M$. \square

Each agent i stores a fixed constraint $h^{[i]} = h_{i0} \cap H_M$, where each h_{i0} is a local polyhedron (e.g., a single half-plane) known only by agent i , and updates two local states, namely $z^{[i]}$, associated to the decision variable z of the MILP, and $B^{[i]}$ being a candidate basis of the problem. At the generic (universal) time instant t , agent i calls CUTORACLE which returns, based on the current state $z^{[i]}(t)$, the MIG cut, $h_{\text{MIG}}(t)$, obtained by MIGORACLE, and a cost-based cut, $h_c(t) = \{c^\top z \geq \sigma^{[i]}(t)\}$, where $\sigma^{[i]}(t)$ must satisfy

$\sigma^{[i]}(t) \geq c^\top z^{[i]}(t)$. Then, agent i solves a local LP in which the common objective function $c^\top z$ is minimized subject to the following constraints: the intersection of its neighbors' candidate bases, $\bigcap_{j \in \mathcal{N}_i(t)} B^{[j]}(t)$, its own candidate basis, $B^{[i]}(t)$, the inequality constraint $h^{[i]}$, MIG cut $h_{\text{MIG}}(t)$, and the cost-based cut $h_c(t)$. This procedure is formalized in the following table.

Distributed Meta-Algorithm DiCUT-MILP

State $(z^{[i]}, B^{[i]})$

Initialization

$$h^{[i]} = h_{i0} \cap H_M$$

$$(z^{[i]}, B^{[i]}) = \text{LPLEXSOLV}(h^{[i]}, c)$$

Evolution

$$(h_{\text{MIG}}(t), h_c(t)) = \text{CUTORACLE}(z^{[i]}(t), B^{[i]}(t), c)$$

$$H_{\text{TMP}}(t) = \left(\bigcap_{j \in \mathcal{N}_i(t)} B^{[j]}(t) \right) \cap B^{[i]}(t) \cap h^{[i]} \cap h_{\text{MIG}}(t) \cap h_c(t)$$

$$(z^{[i]}(t+1), B^{[i]}(t+1)) = \text{LPLEXSOLV}(H_{\text{TMP}}(t), c)$$

Summarizing, at each communication round t , each agent sends to its out-neighbors a candidate basis, consisting of d linear constraints, and receives bases from its in-neighbors. Then, each agent solves a local LP whose number of constraints depends on the dimension d and on the number of in-neighbors. Notice that the distributed meta-algorithm does not require time synchronization: each agent can run its local routine (i.e., generate cuts and solve the local LP) at its own rate, by directly using the available in-neighbor bases.

Next we provide two distributed algorithms. The first one provides an exact solution under the assumption of integer-valued optimal cost. The second algorithm exploits a suitable reformulation and approximation of the centralized problem to compute an ϵ -suboptimal solution.

C. INT-DiCUT-MILP

We introduce a distributed algorithm to solve a MILP in the form (1) under the assumption that the optimal cost is integer. The algorithm was originally introduced in the preliminary conference paper [1]. Following the proposed distributed meta-algorithm, we set $h_{i0} = \{a_i^\top z \leq b_i\}$. Then CUTORACLE consists of an oracle that generates MIG cuts, i.e., $h_{\text{MIG}}(t) = \text{MIGORACLE}(z^{[i]}(t), B^{[i]}(t))$, and another oracle generating cost-based cuts, $h_c(t) = \{c^\top z \geq \sigma^{[i]}(t)\}$, based on a simple ceiling of the current cost, that is, $\sigma^{[i]}(t) = \lceil c^\top z^{[i]}(t) \rceil$. We now state the convergence result for INT-DiCUT-MILP, whose proof is given in the convergence analysis section.

Theorem 3.2 (INT-DiCUT-MILP convergence): Let MILP (1) satisfy Assumptions 2.1 and 3.1, and let the optimal cost be integer-valued. Assume agents communicate according to a jointly strongly connected communication graph, $\mathcal{G}_c(t)$, $t \geq 0$, and run INT-DiCUT-MILP distributed algorithm. Let $J^{[i]}(t) = c^\top z^{[i]}(t)$ be the cost associated to the local candidate lex-optimal solution $z^{[i]}(t)$ of agent $i \in \{1, \dots, N\}$ at time $t \geq 0$. Then, in a finite number of communication rounds the sequences $\{J^{[i]}(t)\}_{t \geq 0}$ and $\{z^{[i]}(t)\}_{t \geq 0}$, $i \in \{1, \dots, N\}$,

converge, respectively, to the (integer-valued) optimal cost and to the lex-optimal solution of problem (1). \square

Consistently with the (centralized) Gomory's Cutting-Plane Algorithm, INT-DiCUT-MILP requires an integer-valued optimal cost. Such assumption is needed in order to guarantee that cost-based cuts (obtained by rounding up the current local cost value $c^\top z^{[i]}(t)$) are valid cutting planes. One way to guarantee integer-valued optimal cost is to take into account MILP instances in which the cost function value depends only on the integer variables and the cost vector has rational components (which, through a suitable scaling, is equivalent to assuming they are integer valued). We point out that for a variety of practical problems, as, for example, scheduling, cutting stock, warehouse location, [33], the optimal cost is integer because all the decision variables are required to be integer and the coefficients of the cost vector c are rational. On the other hand, in many cases of interest, an integer-valued optimal cost cannot be guaranteed a-priori. In the next subsection we remove this assumption and propose a distributed algorithm based on a suitable reformulation and approximation of the centralized problem.

D. ϵ -DiCUT-MILP

Here we propose a distributed algorithm, based on DiCUT-MILP meta algorithm, which computes an ϵ -suboptimal solution for general MILPs. We first consider an equivalent formulation of MILP (1), namely the epigraph form

$$\begin{aligned} & \min_{\rho, z} \rho \\ & \text{subj. to } a_i^\top z \leq b_i, \quad i = 1, \dots, n \\ & \quad c^\top z \leq \rho \\ & \quad \rho \in \mathbb{R}, z \in \mathbb{Z}^{d_z} \times \mathbb{R}^{d_r} \end{aligned} \quad (3)$$

where $\rho \in \mathbb{R}$ is a new decision variable. Problems (1) and (3) are equivalent in the sense that (ρ^*, z^*) is optimal for the epigraph form (3) if and only if z^* is optimal for problem (1) and $\rho^* = c^\top z^*$.

Now, we make a change of variables $\epsilon \rho_I = \rho$, with $\epsilon > 0$, and approximate problem (3) by constraining the new variable ρ_I to be integer. The resulting approximate MILP is

$$\begin{aligned} & \min_{\rho_I, z} \rho_I \\ & \text{subj. to } a_i^\top z \leq b_i, \quad i = 1, \dots, n \\ & \quad c^\top z \leq \epsilon \rho_I \\ & \quad \rho_I \in \mathbb{Z}, z \in \mathbb{Z}^{d_z} \times \mathbb{R}^{d_r}, \end{aligned} \quad (4)$$

where the constant $\epsilon > 0$ is neglected in the objective function since it does not affect the minimization.

With problem (4) at hand, we are ready to introduce ϵ -DiCUT-MILP distributed algorithm. The algorithm is obtained by properly targeting DiCUT-MILP distributed meta-algorithm to problem (4). That is, we set (ρ_I, z) as extended decision variable, e_1 as cost (as opposed to c), and set $h_{i0} = \{a_i z \leq b_i\} \cap \{c^\top z \leq \epsilon \rho_I\}$ as local constraint. The polyhedron induced by the inequality constraints is $P^\epsilon = \{(\rho_I, z) \in \mathbb{R}^{d+1} : a_i^\top z \leq b_i, i = 1, \dots, n, c^\top z \leq \epsilon \rho_I\}$. We denote H_M^ϵ the bounding box associated to the

decision variable (ρ_I, z) . Consistently with the distributed meta-algorithm, each agent generates a MIG cut $h_{\text{MIG}}(t) = \text{MIGORACLE}((\rho_I^{[i]}(t), z^{[i]}(t)), B^{[i]}(t))$ and a cost-based cut $h_c(t) = \{\rho_I \geq \sigma^{[i]}(t)\}$ where $\sigma^{[i]}(t) = \lceil \rho_I^{[i]}(t) \rceil$. Then, it solves a local LP based on the generated cuts and on the neighboring bases. To better highlight the connection with the meta-algorithm, a pseudo-code description of $\epsilon\text{-Dicut-MILP}$ is reported in the following table.

Distributed Algorithm $\epsilon\text{-Dicut-MILP}$

State $((\rho_I^{[i]}, z^{[i]}), B^{[i]})$

Initialization

$$h^{[i]} = h_{i0} \cap H_M$$

$$((\rho_I^{[i]}, z^{[i]}), B^{[i]}) = \text{LPXOLSOLV}(h^{[i]}, e_1)$$

Evolution

$$h_{\text{MIG}}(t) = \text{MIGORACLE}((\rho_I^{[i]}(t), z^{[i]}(t)), B^{[i]}(t))$$

$$h_c(t) = \{\rho_I \geq \lceil \rho_I^{[i]}(t) \rceil\}$$

$$H_{\text{TMP}}(t) = \left(\bigcap_{j \in \mathcal{N}_i(t)} B^{[j]}(t) \right) \cap B^{[i]}(t) \cap h^{[i]} \cap h_{\text{MIG}}(t) \cap h_c(t)$$

$$((\rho_I^{[i]}(t+1), z^{[i]}(t+1)), B^{[i]}(t+1)) = \text{LPXOLSOLV}(H_{\text{TMP}}(t), e_1)$$

The convergence properties of $\epsilon\text{-Dicut-MILP}$ are stated in the next theorem whose proof is given in the analysis section.

Theorem 3.3 ($\epsilon\text{-Dicut-MILP}$ convergence): Let MILP (1) satisfy Assumptions 2.1 and 3.1 with $M \geq \max(-\min_{z \in P} c^\top z / \epsilon, \lceil \max_{z \in P} c^\top z / \epsilon \rceil)$. Assume agents communicate according to a jointly strongly connected communication graph, $\mathcal{G}_c(t)$, $t \geq 0$, and run $\epsilon\text{-Dicut-MILP}$ distributed algorithm. Then, in a finite number of communication rounds the sequences $\{z^{[i]}(t)\}_{t \geq 0}$, $i \in \{1, \dots, N\}$, converge to an ϵ -suboptimal solution of (1). \square

From Theorem 3.3 it follows immediately that if the optimal cost of MILP (1) consists of q decimal digits, then, by setting $\epsilon = 10^{-q}$, agents compute an optimal solution of (1) in a finite number of communication rounds.

Remark 3.4: (Multiple Cuts) Both INT-Dicut-MILP and $\epsilon\text{-Dicut-MILP}$ can be implemented by generating multiple MIG cuts at the generic communication round t . That is, together with the MIG cut with respect to the first non-integer component generated by MIGORACLE , agent i also generates intersection cuts for other non-integer entries. The introduction of multiple cuts can be a useful tool for faster and numerically robust convergence, as shown in the numerical computations. Proofs are provided for a single cut, but the extension to multiple cuts follows by using similar arguments. \square

E. Discussion of Main Algorithm Features

We discuss some interesting features of the proposed algorithms. First, our distributed algorithms only require the communication graph to be jointly strongly connected, and the universal time does not need to be known by the agents. Indeed, agents do not use it in the local updates. This implies, as we will show in the numerical computations, that the algorithms work under asynchronous and unreliable communication networks and, in particular, in networks subject to packet loss.

Second, agents can realize that convergence has occurred, and thus halt the algorithm, in a purely distributed way under slightly stronger assumptions on the graph. We recall that, for a static digraph, the diameter d_G is the maximum distance taken over all the pairs of agents (i, j) , where the distance is defined as the length of the shortest directed path from i to j . Since our distributed algorithms converge in a finite number of communication rounds, it can be shown that, for static communication digraphs, each agent running the algorithm can stop the algorithm if its basis has not changed for $2d_G + 1$ communication rounds, see, e.g., [21]. It is worth noting that, in the initialization step, each agent can compute the graph diameter by a simple flooding algorithm. By similar arguments it can be shown that, if the graph is (uniformly jointly) L -strongly connected, then each agent can stop the algorithm if the value of its basis has not changed after $2LN + 1$ communication rounds. Third, the distributed meta-algorithm (so as the two specific algorithms) involves local computations and communications that depend on the dimension d of the decision variable and on the number of in-neighbors. Indeed, an agent sends to neighbors a candidate basis, which is a collection of d linear constraints, generates cutting planes based on simple local computations, and solves a LP. Thus, the main computational burden for the i -th agent is due to the solution of a LP (solvable in polynomial time) with d variables and $d \times \mathcal{N}_i(t)$ constraints.

Finally, we conclude this discussion by highlighting that the idea of solving a suitable epigraph approximation of the original problem, proposed for $\epsilon\text{-Dicut-MILP}$, can be used also in a centralized setup. This would result into a pure cutting plane algorithm providing an ϵ -suboptimal solution for general MILPs with a computation burden comparable with the Gomory algorithm (which however works only under the assumption of integer-valued optimal cost). On this regard, we point out that other approximate algorithms, as the ones mentioned in Section II-C, require inner, computationally expensive, procedures to find a suboptimal solution.

IV. CONVERGENCE ANALYSIS OF Dicut-MILP ALGORITHMS

In this section we provide a twofold result. First, we provide two technical lemmas that hold for the Dicut-MILP distributed meta-algorithm and that can be used to prove the convergence of a class of distributed optimization algorithms based on this approach. Specifically, we prove that for each agent the local cost and the local state converge in a finite number of communication rounds if the generation of cost-based cuts is assumed to stop in finite time. Then, we prove that under the same condition, consensus among all the agents is attained for the costs and for the candidate lex-optimal solutions. Second, relying on these results, we prove the convergence of INT-Dicut-MILP and $\epsilon\text{-Dicut-MILP}$ distributed algorithms. Namely, for INT-Dicut-MILP we show that agents agree on the lex-optimal solution of MILP (1) under the assumption of integer-valued optimal cost. For $\epsilon\text{-Dicut-MILP}$ we show that they agree on an ϵ -suboptimal solution of MILP (1).

A. Property of MIG Cuts to a Basis

Here, we prove a property, that will be used in our convergence analysis, holding for MIG cuts generated with respect to a basis (rather than with respect to the entire polyhedron as usually done in centralized algorithms).

Lemma 4.1: Let z^{LP} be the lex-optimal solution of a generic LP relaxation of problem (1) with $z_\ell^{\text{LP}} \notin \mathbb{Z}$ for some $\ell \in \{1, \dots, d_Z\}$, and $B^{\text{LP}} = \{A_B z \leq b_B\}$ an associated basis. Let $D(e_\ell, \lfloor z_\ell^{\text{LP}} \rfloor)$ be the split disjunction $\{e_\ell^\top x \leq \lfloor z_\ell^{\text{LP}} \rfloor\} \cup \{e_\ell^\top x \geq \lfloor z_\ell^{\text{LP}} \rfloor + 1\}$, which does not contain z^{LP} . Let $h := \{\alpha^\top z \leq \beta\}$ be the intersection cut to the disjunction $D(e_\ell, \lfloor z_\ell^{\text{LP}} \rfloor)$ and to the basis B^{LP} . Then, the lex-optimal solution, z^B , obtained by minimizing the linear cost $c^\top z$ over $B^{\text{LP}} \cap h$ is such that its ℓ -th component is either $z_\ell^B = \lfloor z_\ell^{\text{LP}} \rfloor$ or $z_\ell^B = \lfloor z_\ell^{\text{LP}} \rfloor + 1$.

Proof: The intersection of the half-planes defining B^{LP} defines a translated cone $C(z^{\text{LP}})$ with apex z^{LP} . Points along an extreme ray² $r^m \in \mathbb{R}^d$, $m \in \{1, \dots, d\}$, associated to $C(z^{\text{LP}})$ are described by $z = z^{\text{LP}} + \mu r^m$, $\mu \in \mathbb{R}$ and $\mu \geq 0$. Thus, vectors $z \in B^{\text{LP}}$ can be described by a positive linear combination of the extreme rays r^m of the basis B^{LP} , i.e., $z = z^{\text{LP}} + \sum_{m=1}^d \mu_m r^m$, $\mu_m \geq 0$. For those r^m such that $e_\ell^\top r^m \neq 0$, let us define $z^m = z^{\text{LP}} + \lambda_m r^m$ with

$$\lambda_m = \begin{cases} \frac{\lfloor z_\ell^{\text{LP}} \rfloor - e_\ell^\top z^{\text{LP}}}{e_\ell^\top r^m} & \text{if } e_\ell^\top r^m < 0 \\ \frac{\lfloor z_\ell^{\text{LP}} \rfloor - e_\ell^\top z^{\text{LP}} + 1}{e_\ell^\top r^m} & \text{if } e_\ell^\top r^m > 0. \end{cases}$$

Then $z^m \in B^{\text{LP}}$ because $\lambda_m \geq 0$. Moreover, z^m is the point obtained as intersection of the ray r^m with the disjunctive hyperplanes $e_\ell^\top z = \lfloor z_\ell^{\text{LP}} \rfloor$ or $e_\ell^\top z = \lfloor z_\ell^{\text{LP}} \rfloor + 1$. In fact, by solving

$$\begin{cases} z = z^{\text{LP}} + \mu r^m \\ e_\ell^\top z = \lfloor z_\ell^{\text{LP}} \rfloor + 1 \end{cases},$$

we have

$$\begin{aligned} e_\ell^\top (z^{\text{LP}} + \mu r^m) &= \lfloor z_\ell^{\text{LP}} \rfloor + 1 \\ \mu e_\ell^\top r^m &= \lfloor z_\ell^{\text{LP}} \rfloor + 1 - e_\ell^\top z^{\text{LP}} \\ \mu &= \frac{\lfloor z_\ell^{\text{LP}} \rfloor + 1 - e_\ell^\top z^{\text{LP}}}{e_\ell^\top r^m}, \end{aligned}$$

which is well defined and nonnegative if $e_\ell^\top r^m > 0$. A similar argument holds for $e_\ell^\top z = \lfloor z_\ell^{\text{LP}} \rfloor$ which gives the condition $e_\ell^\top r^m < 0$. If $e_\ell^\top r^m = 0$, the extreme ray r^m is parallel to the disjunctive hyperplanes $e_\ell^\top z = \lfloor z_\ell^{\text{LP}} \rfloor$ and $e_\ell^\top z = \lfloor z_\ell^{\text{LP}} \rfloor + 1$, [40]. The intersection cut to the split disjunction $D(e_\ell, \lfloor z_\ell^{\text{LP}} \rfloor)$ and the basis B^{LP} , $h = \{\alpha^\top z \leq \beta\}$, is then defined by the hyperplane passing through the intersection points z^m such that $e_\ell^\top r^m \neq 0$ and does not intersect the rays such that $e_\ell^\top r^m = 0$. Now, the lexicographic minimization of the linear cost function over $h \cap B^{\text{LP}}$ returns z^B (cutting off z^{LP}) and the new basis characterized by $d-1$ constraints from B^{LP} and h . The intersection of the half-planes defining the new basis defines a translated cone $C(z^B)$ with apex z^B . Let us consider the extreme ray $r^{\bar{m}}$ defined by the $d-1$ constraints of B^{LP} in the new basis. Now, $r^{\bar{m}}$ cannot be parallel to the intersection cut h otherwise the associated $d-1$ constraints could not

belong to the new basis. Thus, there exists a point $z^{\bar{m}}$ on the ray intersecting one of the two disjunctive hyperplanes. By construction, $z^{\bar{m}}$ belongs also to the intersection cut h and thus it must be the unique intersection point z^B . Therefore, z^B has ℓ -th component equal to $\lfloor z_\ell^{\text{LP}} \rfloor$ or $\lfloor z_\ell^{\text{LP}} \rfloor + 1$. ■

B. Technical Results for the DiCUT-MILP Meta-Algorithm

Lemma 4.2 (Local convergence): Let MILP (1) satisfy Assumptions 2.1 and 3.1. Assume agents run an algorithm based on DiCUT-MILP distributed meta-algorithm such that, at each node, CUTOACLE generates cost-based cuts $\{c^\top z \geq \sigma^{[i]}(t)\}$ with $\sigma^{[i]}(t)$ taking values in a finite set. Then, in a finite number of communication rounds, for all $i \in \{1, \dots, N\}$,

- i) the sequence $\{J^{[i]}(t)\}_{t \geq 0}$ converges to a constant value $\bar{J}^{[i]}$, and
- ii) the sequence $\{z^{[i]}(t)\}_{t \geq 0}$ converges to a feasible $\bar{z}^{[i]}$ (with $\bar{z}_1^{[i]}, \dots, \bar{z}_{d_Z}^{[i]}$ integer).

Proof: To prove that sequence $\{J^{[i]}(t)\}_{t \geq 0}$ converges to a constant $\bar{J}^{[i]}$ in a finite number of communication rounds we proceed in four steps.

First, the sequence $\{J^{[i]}(t)\}_{t \geq 0}$ is monotonically non-decreasing. Indeed, to compute $B^{[i]}(t+1)$, the i -th agent minimizes the common objective function subject to the constraint set $H_{\text{TMP}}(t)$ which, by construction, is a subset of the basis $B^{[i]}(t)$. Second, the sequence is bounded from above by the optimal cost J^* to (1) (since $H_{\text{TMP}}(t)$ always includes $\text{conv}(P_I)$). Thus, $J^{[i]}(t)$ converges to some $\bar{J}^{[i]} \leq J^*$. To conclude the proof of statement i), we recall that at each t , $h_c(t) = \{c^\top z \geq \sigma^{[i]}(t)\}$, with $\sigma^{[i]}(t) \geq J^{[i]}(t)$ and by assumption $\sigma^{[i]}(t)$ takes value in a finite set.

To prove the second statement, first notice that we have just proved that there exists a time t_0 such that $J^{[i]}(t) = \bar{J}^{[i]}, \forall t \geq t_0$. Now let us consider the sequence of the first component of the state associated to the integer decision variable, i.e., $\{z_1^{[i]}(t)\}_{t \geq t_0}$. The sequence $\{z_1^{[i]}(t)\}_{t \geq t_0}$ is non-decreasing because the local cost value is constant after t_0 , $H_{\text{TMP}}(t)$ is a subset of $B^{[i]}(t)$, $\forall t \geq t_0$, and the sequence is constructed by taking into account the lex-optimal solution of the local problem. Moreover, it is upper bounded by M , and therewith convergent with limit $\bar{z}_1^{[i]}$. So, there exists a time $t_1 \geq t_0$ such that $\lceil \bar{z}_1^{[i]} \rceil - 1 < z_1^{[i]}(t) \leq \lceil \bar{z}_1^{[i]} \rceil, \forall t \geq t_1$. Following the evolution of the meta-algorithm, to compute $z_1^{[i]}(t_1+1)$, first agent i generates, through CUTOACLE, a MIG cut $h_{\text{MIG}}(t_1)$ to the split disjunction $D(e_1, \lceil \bar{z}_1^{[i]} \rceil - 1)$ and to the current basis $B^{[i]}(t_1)$. Then it collects the constraints from its neighbors, builds up $H_{\text{TMP}}(t_1)$, and calls LPLEXSOLV. This returns a new lex-optimal solution $z_1^{[i]}(t_1+1)$, which is greater than or equal to the solution obtained by minimizing over $B^{[i]}(t_1) \cap h_{\text{MIG}}(t_1)$, since $H_{\text{TMP}}(t_1)$ is a subset of it. By Lemma 4.1, minimizing the linear cost function over the set of the current basis and the intersection cut is such that the first component is either $\lceil \bar{z}_1^{[i]} \rceil - 1$ or $\lceil \bar{z}_1^{[i]} \rceil$. But, being $\lceil \bar{z}_1^{[i]} \rceil - 1 < z_1^{[i]}(t)$ for all $t \geq t_1$, it must hold $z_1^{[i]}(t_1+1) = \lceil \bar{z}_1^{[i]} \rceil$ and, thus, $z_1^{[i]}(t) = \lceil \bar{z}_1^{[i]} \rceil$ for all $t > t_1 + 1$. Therefore, we have shown that the sequence $\{z_1^{[i]}(t)\}_{t \geq t_0}$ converges to $\bar{z}_1^{[i]} = \lceil \bar{z}_1^{[i]} \rceil$ in a finite number of communication rounds. The same argument

²A nonzero vector r of a polyhedral cone C is called an extreme ray if there are $d-1$ linearly independent constraints that are active at r , [39].

can now be applied to the remaining integer components of the decision vector, $z_2^{[i]}, \dots, z_{d_z}^{[i]}$. So there exists a time t_{d_z} such that $(z_1^{[i]}(t), \dots, z_{d_z}^{[i]}(t)) = (\bar{z}_1^{[i]}, \dots, \bar{z}_{d_z}^{[i]})$, $\forall t \geq t_{d_z}$, and thus agent i will not generate cutting planes anymore. Let T_{d_z} be such that $(z_1^{[i]}(t), \dots, z_{d_z}^{[i]}(t)) \in \mathbb{Z}^{d_z}$, $\forall t \geq T_{d_z}$, $\forall i \in \{1, \dots, N\}$. For $t \geq T_{d_z}$, no cutting planes will be generated in the network. This means that the number of possible different bases that i -th agent can receive from its neighbors is finite (specifically, the combination of all the constraints in the network at T_{d_z}). Therefore, since the lexicographic cost is nondecreasing and due to the finite number of possible bases, also the non-integer variables $(z_{d_z+1}^{[i]}(t), \dots, z_d^{[i]}(t))$ will converge in a finite number of communication rounds, thus concluding the proof. ■

Next we prove that agents reach consensus on (a common) cost value and solution estimate in a finite number of communication rounds. The proof follows similar arguments as in [21] but we report all the steps for the sake of completeness.

Lemma 4.3 (Consensus): Assume that the sequences $\{J^{[i]}(t)\}_{t \geq 0}$ and $\{z^{[i]}(t)\}_{t \geq 0}$ defined as in Lemma 4.2 converge to a constant value $\bar{J}^{[i]}$ and a feasible $\bar{z}^{[i]}$, respectively, $\forall i \in \{1, \dots, N\}$. Assume the communication network, $\mathcal{G}_c(t)$, is jointly strongly connected. Then, $\bar{J}^{[i]} = \bar{J}^{[j]}$ and $\bar{z}^{[i]} = \bar{z}^{[j]}$ for all $i, j \in \{1, \dots, N\}$.

Proof: We start by proving the consensus in cost. Suppose, by contradiction, that two cost sequences, associated with two agents, say i and j , converge to two different values $\bar{J}^{[i]}$ and $\bar{J}^{[j]}$, respectively. Let, without loss of generality, $\bar{J}^{[j]} > \bar{J}^{[i]}$. In Lemma 4.2 we have shown that both sequences, $\{J^{[i]}(t)\}_{t \geq 0}$ and $\{J^{[j]}(t)\}_{t \geq 0}$, are monotonically non-decreasing and convergent in a finite number of communication rounds. Since $\bar{J}^{[j]} - \bar{J}^{[i]} > 0$, there must exist $T_{\delta_{ij}} > 0$ such that $\bar{J}^{[j]} \geq J^{[j]}(t) > \bar{J}^{[i]} \geq J^{[i]}(t)$, $\forall t > T_{\delta_{ij}}$. Now since the communication graph is jointly strongly connected, for each time $t \geq 0$ and each pair of agents (i, j) , there exists a sequence of nodes $\{l_1, \dots, l_k\}$ and an increasing sequence of time instants $\{t_1, \dots, t_{l_k+1}\}$, with $t \leq t_{l_1} < \dots < t_{l_k+1}$, such that the directed edges $\{(j, l_1), (l_1, l_2), \dots, (l_k, i)\}$ belong to the digraph at times $\{t_1, \dots, t_{l_k+1}\}$, [21]. To compute $B^{[l_1]}(t_{l_1} + 1)$, agent l_1 minimizes the common objective function subject to the constraint set $H_{\text{TMP}}(t_{l_1})$ which, by construction, is a subset of the basis $B^{[j]}(t_{l_1})$. Therefore, $J^{[l_1]}(t_{l_1} + 1) \geq J^{[j]}(t_{l_1})$. Iterating, we have $J^{[i]}(t_{l_k+1} + 1) \geq J^{[j]}(t_{l_1})$. Therefore, for each $t > T_{\delta_{ij}}$ there exists $\tau > 0$ such that $J^{[i]}(t + \tau) \geq J^{[j]}(t)$, which leads to a contradiction, thus proving that $\bar{J}^{[1]} = \dots = \bar{J}^{[N]}$.

To prove that $\bar{z}^{[i]} = \bar{z}^{[j]} \forall i, j \in \{1, \dots, N\}$, first recall that each sequence $\{z^{[i]}(t)\}_{t \geq 0}$ converges to $\bar{z}^{[i]} \forall i \in \{1, \dots, N\}$ in a finite number of communication rounds. Let us suppose, by contradiction, that agents do not reach consensus on the decision variable, then there exist two agents, say i and j , such that $\bar{z}^i >_{\text{lex}} \bar{z}^j$. Using again the assumption that the graph is jointly strongly connected, there exists a directed path $\{(i, l_1), (l_1, l_2), \dots, (l_k, j)\}$ at times $\{t_1, \dots, t_{l_k+1}\}$. Since consensus has been reached on the cost, it must hold (by applying to the lexicographic ordering the same argument used for the cost) that $\bar{z}^j = z^{[j]}(t_{l_k+1} + 1) \geq_{\text{lex}} z^{[l_k]}(t_{l_k+1}) \geq_{\text{lex}}$

$\dots \geq_{\text{lex}} z^{[i]}(t) = \bar{z}^i$ which leads to a contradiction, thus concluding the proof. ■

C. Proof of Theorem 3.2

We apply Lemma 4.2 to MILP (1). Assumptions in Lemma 4.2 are satisfied. Indeed, at each node, the number of cost-based cuts, $h_c = \{c^\top z \geq \sigma^{[i]}(t)\}$, that can be generated along the algorithm evolution is finite, because $\sigma^{[i]}(t) = \lceil c^\top z^{[i]}(t) \rceil$ and $c^\top z^{[i]}(0) \leq c^\top z^{[i]}(t) \leq J^*$ for all t , with J^* being the (integer-valued) optimal cost to (1). Therefore, by Lemma 4.2, the sequences $\{J^{[i]}(t)\}_{t \geq 0}$ and $\{z^{[i]}(t)\}_{t \geq 0}$ at each node i converge in a finite number of communication rounds. Now, since the communication graph is jointly strongly connected, by Lemma 4.3 the sequences converge in a finite number of communication rounds to a common limit cost value \bar{J} and a common limit point \bar{z} respectively.

To prove that $\bar{J} = J^*$, we start by highlighting two facts. First, $J^{[i]}(t) \leq J^*$, $\forall t \geq 0$, and for all $i \in \{1, \dots, N\}$, because every agent minimizes the cost function over $H_{\text{TMP}}(t)$, which by construction (it is the intersection of the collected bases and the local generated cuts) always contains $\text{conv}(P_I)$. Thus, $\bar{J} \leq J^*$. Second, there exists \bar{t} such that $z^{[i]}(t) = \bar{z}$ for all $t \geq \bar{t}$ and $\forall i \in \{1, \dots, N\}$. This means that \bar{z} satisfies the local constraints $h^{[i]}$ for all $i \in \{1, \dots, N\}$ and, therefore, $\bar{z} \in \left(\bigcap_{i=1}^N h^{[i]} \right) = P$. Moreover, since $\bar{z} = (\bar{x}, \bar{y}) \in P$ is such that $\bar{x} \in \mathbb{Z}^{d_z}$, it holds $\bar{z} \in \text{conv}(P_I)$ by construction. Thus, $J(\bar{z}) \geq \min_{z \in \text{conv}(P_I)} J(z) = J^*$, giving the cost optimality. Finally, being \bar{z} feasible and cost optimal, it is the unique lex-optimal solution (i.e., $\bar{z} = z^*$).

D. Proof of Theorem 3.3

In this subsection, we prove Theorem 3.3, i.e., ϵ -Dicut-MILP computes an ϵ -suboptimal solution of MILP (1). Before proving the convergence result, we need the following lemma.

Lemma 4.4 (Approximate solution): Suppose MILP (1) satisfies Assumption 2.1. Then for any given $\epsilon > 0$ an optimal solution $(\rho_I^\epsilon, z^\epsilon)$ to MILP (4) exists and z^ϵ is an ϵ -suboptimal (feasible) solution to MILP (1).

Proof: By Assumption 2.1, there exists $z^* \in \text{conv}(P_I)$ such that $c^\top z^* \leq c^\top z, \forall z \in \text{conv}(P_I)$. If $c^\top z^*/\epsilon \in \mathbb{Z}$, then $(\rho_I^\epsilon, z^\epsilon) = (c^\top z^*/\epsilon, z^*)$ is a feasible, optimal solution of (4). Any optimal solution $(\rho_I^\epsilon, z^\epsilon)$ of (4) has cost $c^\top z^*/\epsilon$ and, by construction, $z^\epsilon \in \text{conv}(P_I)$. Thus, $c^\top z^\epsilon/\epsilon \leq c^\top z^*/\epsilon = \rho_I^\epsilon$ and, being $z^\epsilon \in \text{conv}(P_I)$, $c^\top z^\epsilon \geq c^\top z^*$, so that z^ϵ is also optimal for (1). Now suppose $c^\top z^*/\epsilon \notin \mathbb{Z}$. We first show that the pair $\rho_I^\epsilon = \lceil c^\top z^*/\epsilon \rceil$, $z^\epsilon = z^*$ is an optimal solution of (4). Indeed, z^* is feasible for (4), i.e., $z^* \in \text{conv}(P_I)$, and $c^\top z^* \leq \epsilon \lceil c^\top z^*/\epsilon \rceil = \epsilon \rho_I^\epsilon$ (simply because $c^\top z^*/\epsilon \leq \lceil c^\top z^*/\epsilon \rceil$). Moreover, there does not exist any $z \in \text{conv}(P_I)$ such that $c^\top z = \epsilon \lfloor c^\top z^*/\epsilon \rfloor$. In fact, by the principle of optimality, we have $c^\top z \geq c^\top z^*, \forall z \in \text{conv}(P_I)$, and $c^\top z^* > \epsilon \lfloor c^\top z^*/\epsilon \rfloor$ (we have already handled the case of integer-valued optimal cost). This proves the existence of an optimal solution of MILP (4). Now we simply notice that any

$(\rho_I^\epsilon, z^\epsilon)$ optimal solution of (4) is such that $z^\epsilon \in \text{conv}(P_I)$ and $c^\top z^\epsilon \leq \epsilon \rho_I^\epsilon = \epsilon \lceil c^\top z^* / \epsilon \rceil$. Moreover, $c^\top z^\epsilon \geq c^\top z^*$ because $z^\epsilon \in \text{conv}(P_I)$, so that $c^\top z^* \leq c^\top z^\epsilon \leq \epsilon \lceil c^\top z^* / \epsilon \rceil$. By subtracting $c^\top z^*$, we have

$$0 \leq c^\top z^\epsilon - c^\top z^* \leq \epsilon \lceil c^\top z^* / \epsilon \rceil - c^\top z^* < \epsilon,$$

thus concluding the proof. \blacksquare

We are now ready to prove that ϵ -DiCUT-MILP converges in a finite number of communication rounds to an ϵ -suboptimal solution of MILP (1).

Proof of Theorem 3.3: First, note that we can prove the theorem by replacing the feasible set of MILP (4) with $P^\epsilon = \{(\rho_I, z) \in \mathbb{R}^{d+1} : a_i^\top z \leq b_i, i = 1, \dots, n, c^\top z \leq \epsilon \rho_I, \rho_I \leq M_\rho\}$, where $M_\rho \geq \lceil \max_{z \in P} c^\top z / \epsilon \rceil$. Here we are slightly abusing notation, since we denote by P^ϵ the bounded version of the polyhedron. Using P^ϵ it can be shown that the assumptions of Theorem 3.2 hold for MILP (4). Moreover, the optimal cost of MILP (4) is integer (indeed, the cost depends only on the integer variable ρ_I) and the communication graph is jointly strongly connected. Thus, the sequences $\{J^{[i]}(t)\}_{t \geq 0}$ and $\{(\rho_I^{[i]}(t), z^{[i]}(t))\}_{t \geq 0} \ i \in \{1, \dots, N\}$, converge, respectively, to the optimal cost, J^ϵ , and the lex-optimal solution, $(\rho_I^\epsilon, z^\epsilon)$, of MILP (4) in a finite number of communication rounds. Finally, by Lemma 4.4, the solution z^ϵ is an ϵ -suboptimal solution of MILP (1). \blacksquare

V. NUMERICAL COMPUTATIONS

To corroborate the theoretical analysis and better highlight some appealing features of our algorithms, we provide numerical computations. We concentrate on ϵ -DiCUT-MILP since it does not require integer-valued optimal cost. First, we consider randomly generated MILP instances and we perform a numerical Monte Carlo analysis of the algorithm convergence while varying network size and tolerance ϵ . Second, we consider a multi-agent multi-task assignment problem that is relevant for cooperative robotics. For this scenario we test our algorithm in an unreliable network due to packet loss.

A. Scalability and Convergence Rate for Randomly Generated MILPs

We first perform a numerical Monte Carlo analysis to study the number of communication rounds for convergence while varying the network size. For each network size, we randomly generate 50 MILP instances with fixed dimension, $d = 10$ and $d_Z = 3$, and run ϵ -DiCUT-MILP with $\epsilon = 0.1$ and a bounding box H_M with $M = 100$. The inequality constraints of the MILP instances are randomly generated as follows. The vectors a_i are drawn from the standard Gaussian distribution. The term b_i is uniformly randomly generated in $[0, 50]$. The cost vector c is obtained as $c = a_i^\top \hat{c}$ where \hat{c} is uniformly randomly generated in $[0, 1]^n$. The problems generated according to this model are feasible. The LP relaxation of such a MILP is one of the models proposed in [41]. We first choose a cyclic digraph for which the diameter, d_G , is proportional to the number of agents, specifically $d_G = N - 1$, and consider the following cases: number of agents equal to 16, 32, 64 and 128. The results are shown in Figure 3a. The red central line of each




Fig. 3. Communication rounds for convergence while varying the number of agents for (a) a cyclic digraph with $d_G = N - 1$ and (b) an Erdős-Rényi random graph with diameter $d_G = 8$. Each box plot shows the minimum and maximum communication rounds (whiskers), 25% and 75% percentiles (lower and upper box edges), median (red line), and outliers (red crosses).

box shows the median value of the communication rounds for the 50 random MILPs. The lower and upper edges of the blue box represent the 25-th and 75-th percentiles. Some outliers³ (red crosses) can be observed. This is not surprising considered that MILPs are \mathcal{NP} -hard. We point out that the number of communication rounds needed for convergence grows linearly with the network size (i.e., number of agents and diameter). Then, we consider Erdős-Rényi static digraphs with diameter fixed $d_G = 8$, and a growing number of nodes (and so agents), namely 25, 50, 75 and 100. The results are shown in Figure 3b. Here we highlight that the median value of communication rounds needed for convergence is between 177 and 190, that is the completion time of the algorithm is roughly constant with respect to the number of agents.

Second, we perform a numerical Monte Carlo analysis to study the behavior of ϵ -DiCUT-MILP while varying ϵ . We randomly generate 50 MILP instances with $d = 10$, $d_Z = 3$ and 256 constraints, as before. The constraints are distributed among $N = 64$ agents, so that each one knows 4 constraints. Agents communicate according to an Erdős-Rényi random graph with $d_G = 7$. Each instance is solved with $\epsilon = 0.05, 0.1, 0.5, 1$, and the results are shown in Figure 4. For each communication round t , we plot the mean value (over instances and agents), $\text{Avg}(|c^\top z^* - \epsilon \rho_I^{[i]}(t)|)$, of the distance $|c^\top z^* - \epsilon \rho_I^{[i]}(t)|$ of each instance from the optimal cost. Different curves are associated to different values of ϵ . As expected, by decreasing ϵ , the cost at convergence is closer to the optimal one, but a larger number of communication rounds is needed for convergence. To speed-up the convergence, a “multi-resolution” strategy may be implemented, in which ϵ is reduced at each run and the previous computed solution is used to initialize the new run. Notice that, if agents can apply one of the distributed stopping criteria discussed in the previous section, e.g., if the communication graph is static, this scheme can be implemented in a purely distributed way. Finally, we point out that the algorithm exhibits a convergence rate behavior similar to centralized cutting-plane algorithms. That is, in the first iterations the cost soars to the optimal cost value. As the algorithm evolves, one can observe the typical tailing-off

³Outliers are evaluated as follows. Let Q_1 and Q_3 be the 25-th and 75-th percentile of the samples, respectively. A sample is an outlier if it is greater than $Q_3 + 1.5(Q_3 - Q_1)$ or less than $Q_1 - 1.5(Q_3 - Q_1)$.




Fig. 4. Convergence behavior of ϵ -DiCUT-MILP, with different ϵ values. For each communication round t , the figure shows the value of $|c^\top z^* - \epsilon \rho_I^{[i]}(t)|$ averaged over all the agents and all the random instances.

effect characterizing cutting planes algorithms, which slows down the convergence rate. How to combine cutting planes with other schemes, to avoid this tailing-off, is definitely an interesting research direction.

B. Distributed Multi-Agent Multi-Task Assignment under Unreliable Network

In this section we apply ϵ -DiCUT-MILP to a Multi-Agent Multi-Task Assignment problem for robotic networks. Here we use the variant of ϵ -DiCUT-MILP described in Remark 3.4 in which multiple cuts are generated.

We consider a group of heterogeneous mobile robots that need to accomplish a set of tasks. Each task must be performed at a given target location and each vehicle has the capability to execute only some of the tasks. Robots move in a constrained space including obstacles, which (for simplicity) are modeled as rectangles. Given the capability constraints and forbidden areas, the goal is to assign the tasks among the vehicles in order to minimize the mission completion time which is defined as the time needed for the last vehicle to finish its mission, see, e.g., [4]. We assume that valid paths, from vehicles to target locations, are available together with the corresponding finishing times. This could model, e.g., a structured environment, as in factories or warehouses, in which vehicles follow pre-defined paths on the floor of the operating area. Thus, we setup the following multi-agent multi-task assignment problem.

Given N_θ target locations, $\{T_1, \dots, T_{N_\theta}\}$, N_ν vehicles, $\{V_1, \dots, V_{N_\nu}\}$, and N_γ paths, we formulate the multi-agent multi-task assignment problem as the following MILP:

$$\begin{aligned}
 & \min_{x,y} y \\
 & \text{subj. to } p_\theta^\top x \leq -w_\theta, & \forall \theta \in \{1, \dots, N_\theta\} \\
 & q_\nu^\top x \leq 1, & \forall \nu \in \{1, \dots, N_\nu\} \\
 & r_\gamma^\top x \leq y, & \forall \gamma \in \{1, \dots, N_\gamma\}
 \end{aligned} \quad (5)$$

where $x \in \mathbb{Z}^{N_\gamma}$ is a vector of binary decision variables associated to the paths (x_ℓ is equal to one if the ℓ -th path is selected, and 0 otherwise), while $y \in \mathbb{R}$ is a continuous

decision variable associated to the mission completion time. The vector $p_\theta \in \mathbb{R}^{N_\gamma}$ is a vector whose ℓ -th entry is 1 if target location T_θ is visited by path ℓ and 0 otherwise. The vector $q_\nu \in \mathbb{R}^{N_\gamma}$ is a vector whose ℓ -th entry is 1 if path ℓ is assigned to vehicle V_ν and 0 otherwise. Each vector $r_\gamma \in \mathbb{R}^{N_\gamma}$ has the form $r_\gamma^\top = [0, \dots, 0, \tau_\ell, 0, \dots, 0]$, where τ_ℓ is the completion time for the ℓ -th path. The first set of constraints enforces that each target location is visited at least w_θ times, the second one prevents more than one path being assigned to each vehicle, and the third one forces y to be the overall finishing time, i.e., $y = \max_\ell \tau_\ell$. It is worth noting that an optimal solution of (5) cannot be obtained as the solution of its LP relaxation, i.e., $x \in \mathbb{R}^{N_\gamma}$ and $0 \leq x_\ell \leq 1, \forall \ell \in \{1, \dots, N_\gamma\}$, due to the presence of the capability constraints.

We consider N agents, $\{A_1, \dots, A_N\}$, randomly located in the operating region and communicating according to a jointly strongly connected communication graph. The graph is built by considering a proximity criterion, i.e., two agents are connected if their distance is less than a threshold.

Each agent only knows the constraints in problem (5) associated to paths traversing a neighboring area. For example, in Figure 5a, agent A_i (blue circle) only knows the paths traversing a circle area of radius 14m centered at its position, and communicates with other agents in this circle. In this framework, a robotic vehicle, e.g., V_i in Figure 5a, is also an agent participating to the computation.

In order to show the robustness of ϵ -DiCUT-MILP against failures in the communication network, we introduce a probability of packet loss for each edge in the fixed communication graph. Specifically, at each communication round, an agent i does not receive a message from a neighboring agent j if a random variable in $[0, 1]$ is smaller than a fixed threshold. The threshold defines the percentage of packet loss for edge (j, i) . We consider 0% (ideal case, no packet loss), 10%, 30%, 50% and 70% of packet loss. For each packet loss percentage, we run ϵ -DiCUT-MILP on an instance of problem (5) with $\epsilon = 0.1$, $N_\theta = 32$, $N_\nu = 10$, $N_\gamma = 71$, and $N = 30$. The solution is depicted in Figure 5a (colored lines). The convergence results are shown in Figure 5b in which we plot the difference between $\max_{i \in \{1, \dots, N\}} \rho_I^{[i]}(t)$ and the solution ρ_I^* of MILP (4) (computed by using a solver in YALMIP [42]). For the ideal case (no packet loss), agents reach consensus on the optimal cost after 19 communication rounds. When increasing the percentage of packet loss, as expected, the number of communication rounds increases, see Figure 5b. Notably, as from the theory, the algorithm converges also for very high percentages of packet loss as, e.g., 70%.

VI. CONCLUSION

In this paper, we proposed an algorithmic framework, with two specific algorithms, to solve Mixed-Integer Linear Programs over networks. In the proposed distributed setup, the constraints of the MILP are assigned to a network of agents. The agents have a limited amount of memory and computation capabilities and are able to communicate with neighboring agents. Following the idea of centralized cutting-plane methods for MILPs, each agent solves local LP relaxations of the





Fig. 5. In (a) Multi-Agent Multi-Task assignment solution: 10 vehicles (black triangles), 32 tasks (green diamonds), 10 optimal paths (colored lines) and 30 agents (blue circles) connected by a proximity graph (blue dotted lines). In (b) $\max_{i \in \{1, \dots, N\}} (\rho_I^{[i]}(t) - \rho_I^*)$ for different percentages of packet loss.

global problem, generates cutting planes, and exchanges active constraints (a candidate basis) with neighbors. We proved that agents running the first algorithm, INT-DiCUT-MILP, reach consensus on the lex-optimal solution of the MILP, under the assumption of integer-valued optimal cost, in a finite number of communication rounds. Removing the assumption of integer-valued optimal cost, we proved that agents running the algorithm ϵ -DiCUT-MILP reach consensus on a suboptimal solution (up to a given tolerance ϵ) of the MILP, in a finite number of communication rounds. Both algorithms involve low-cost local computations and work under asynchronous, unreliable, and directed networks. Finally, we performed a set of numerical computations suggesting that the completion time of ϵ -DiCUT-MILP scales nicely with the diameter of the communication graph. We also tested the algorithm on a multi-agent multi-task assignment setup in unreliable networks. Future investigations may include the solution of MILPs with a large number of decision variables and the combination of pure cutting plane methods with other tools, e.g. branch and bound, as in centralized schemes.

ACKNOWLEDGMENT

The authors would like to thank Andrea Camisa for the deep discussions and useful suggestions.

REFERENCES

- [1] A. Testa, A. Rucco, and G. Notarstefano, "A finite-time cutting plane algorithm for distributed mixed integer linear programming," in *IEEE Conf. on Decision and Control*, 2017, pp. 3847–3852.
- [2] A. Bemporad and M. Morari, "Control of systems integrating logic, dynamics, and constraints," *Automatica*, vol. 35, no. 3, pp. 407–427, 1999.
- [3] A. Richards, T. Schouwenaars, J. P. How, and E. Feron, "Spacecraft trajectory planning with avoidance constraints using mixed-integer linear programming," *Journal of Guidance, Control, and Dynamics*, vol. 25, no. 4, pp. 755–764, 2002.
- [4] J. Bellingham, M. Tillerson, A. Richards, and J. P. How, "Multi-task allocation and path planning for cooperating UAVs," in *Cooperative control: models, applications and algorithms*. Springer, 2003, pp. 23–41.
- [5] R. E. Gomory, "An algorithm for integer solutions to linear programs," *Recent advances in mathematical programming*, vol. 64, pp. 260–302, 1963.
- [6] A. Bemporad, F. Borrelli, and M. Morari, "Piecewise linear optimal controllers for hybrid systems," in *American Control Conference*. IEEE, 2000, pp. 1190–1194.
- [7] A. Bemporad, F. Borrelli, M. Morari *et al.*, "Model predictive control based on linear programming - The explicit solution," *IEEE Transactions on Automatic Control*, vol. 47, no. 12, pp. 1974–1985, 2002.
- [8] V. Dua and E. N. Pistikopoulos, "An algorithm for the solution of multiparametric mixed integer linear programming problems," *Annals of operations research*, vol. 99, no. 1-4, pp. 123–139, 2000.
- [9] D. Axehill, T. Besselmann, D. M. Raimondo, and M. Morari, "A parametric branch and bound approach to suboptimal explicit hybrid MPC," *Automatica*, vol. 50, no. 1, pp. 240–246, 2014.
- [10] F. Borrelli, D. Subramanian, A. U. Raghunathan, and L. T. Biegler, "MILP and NLP techniques for centralized trajectory planning of multiple unmanned air vehicles," in *American Control Conference*. IEEE, 2006, pp. 5764–5769.
- [11] S. Chopra and M. Egerstedt, "Spatio-temporal multi-robot routing," *Automatica*, vol. 60, pp. 173–181, 2015.
- [12] R. Takapoui, N. Moehle, S. Boyd, and A. Bemporad, "A simple effective heuristic for embedded mixed-integer quadratic programming," *International Journal of Control*, pp. 1–11, 2017.
- [13] S.-J. Kim and G. B. Giannakis, "Scalable and robust demand response with mixed-integer constraints," *IEEE Transactions on Smart Grid*, vol. 4, no. 4, pp. 2089–2099, 2013.
- [14] 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*, vol. 67, no. 5, pp. 144–156, 2016.
- [15] A. Falsone, K. Margellos, and M. Prandini, "A decentralized approach to multi-agent MILPs: finite-time feasibility and performance guarantees," *arXiv preprint arXiv:1706.08788*, 2017.
- [16] D. Richert and J. Cortes, "Distributed linear programming with event-triggered communication," *SIAM Journal on Control and Optimization*, vol. 54, no. 3, pp. 1769–1797, 2016.
- [17] D. Richert and J. Cortés, "Distributed bargaining in dyadic-exchange networks," *IEEE Transactions on Control of Network Systems*, vol. 3, no. 3, pp. 310–321, 2016.
- [18] C. Fischione, M. D'Angelo, and M. Butussi, "Utility maximization via power and rate allocation with outage constraints in nakagami-lognormal channels," *IEEE Transactions on Wireless Communications*, vol. 10, no. 4, pp. 1108–1120, 2011.
- [19] E. Wei, A. Ozdaglar, and A. Jadbabaie, "A distributed newton method for network utility maximization, part I: Algorithm," *IEEE Transactions on Automatic Control*, vol. 58, no. 9, pp. 2162–2175, 2013.
- [20] —, "A distributed newton method for network utility maximization, part II: Convergence," *IEEE Transactions on Automatic Control*, vol. 58, no. 9, pp. 2176–2188, 2013.
- [21] G. Notarstefano and F. Bullo, "Distributed abstract optimization via constraints consensus: Theory and applications," *IEEE Transactions on Automatic Control*, vol. 56, no. 10, pp. 2247–2261, 2011.
- [22] M. Bürger, G. Notarstefano, F. Bullo, and F. Allgöwer, "A distributed simplex algorithm for degenerate linear programs and multi-agent assignments," *Automatica*, vol. 48, no. 9, pp. 2298–2304, 2012.

- [23] S. Chopra, G. Notarstefano, M. Rice, and M. Egerstedt, "Distributed version of the hungarian method for a multirobot assignment," *IEEE Transactions on Robotics*, 2017.
- [24] S. Karaman and G. Inalhan, "Large-scale task/target assignment for UAV fleets using a distributed branch and price optimization scheme," *IFAC Proceedings Volumes*, vol. 41, no. 2, pp. 13 310–13 317, 2008.
- [25] H.-L. Choi, L. Brunet, and J. P. How, "Consensus-based decentralized auctions for robust task allocation," *IEEE Transactions on Robotics*, vol. 25, no. 4, pp. 912–926, 2009.
- [26] M. Franceschelli, D. Rosa, C. Seatzu, and F. Bullo, "Gossip algorithms for heterogeneous multi-vehicle routing problems," *Nonlinear Analysis: Hybrid Systems*, vol. 10, pp. 156–174, 2013.
- [27] Y. Kuwata and J. P. How, "Cooperative distributed robust trajectory optimization using receding horizon MILP," *IEEE Transactions on Control Systems Technology*, vol. 19, no. 2, pp. 423–431, 2011.
- [28] 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*, vol. 2, no. 4, pp. 563–568, 2018.
- [29] M. Bürger, G. Notarstefano, and F. Allgower, "A polyhedral approximation framework for convex and robust distributed optimization," *IEEE Transactions on Automatic Control*, vol. 59, no. 2, pp. 384–395, 2014.
- [30] C. N. Jones, E. C. Kerrigan, and J. M. Maciejowski, "Lexicographic perturbation for multiparametric linear programming with applications to control," *Automatica*, vol. 43, no. 10, pp. 1808–1816, 2007.
- [31] J. E. Kelley, Jr, "The cutting-plane method for solving convex programs," *Journal of the society for Industrial and Applied Mathematics*, vol. 8, no. 4, pp. 703–712, 1960.
- [32] A. Schrijver, *Theory of linear and integer programming*. John Wiley & Sons, 1998.
- [33] M. Jünger, T. M. Lieblich, 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.
- [34] R. E. Gomory, "An algorithm for the mixed integer problem." *RM-2597, The RAND Corporation*, 1960.
- [35] W. Cook, R. Kannan, and A. Schrijver, "Chvátal closures for mixed integer programming problems," *Mathematical Programming*, vol. 47, no. 1, pp. 155–174, 1990.
- [36] E. Balas, "Intersection cuts: A new type of cutting planes for integer programming," *Operations Research*, vol. 19, no. 1, pp. 19–39, 1971.
- [37] M. Jörg, "k-disjunctive cuts and cutting plane algorithms for general mixed integer linear programs," Ph.D. dissertation, Universität München, 2008.
- [38] J. H. Owen and S. Mehrotra, "A disjunctive cutting plane procedure for general mixed-integer linear programs," *Mathematical Programming*, vol. 89, no. 3, pp. 437–448, 2001.
- [39] D. Bertsimas and J. N. Tsitsiklis, *Introduction to linear optimization*. Athena Scientific Belmont, MA, 1997, vol. 6.
- [40] K. Andersen, G. Cornuéjols, and Y. Li, "Split closure and intersection cuts," *Mathematical programming*, vol. 102, no. 3, pp. 457–493, 2005.
- [41] M. J. Todd, "Probabilistic models for linear programming," *Mathematics of Operations Research*, vol. 16, no. 4, pp. 671–693, 1991.
- [42] J. Lofberg, "YALMIP: A toolbox for modeling and optimization in MATLAB," in *IEEE International Symposium on Computer Aided Control Systems Design*, 2004, pp. 284–289.



Andrea Testa received the Laurea degree "summa cum laude" in Computer Engineering from the Università del Salento, Lecce, Italy in 2016.

He is a Ph.D. student in Engineering of Complex Systems at the Università del Salento, Lecce, Italy since November 2016. He was a visiting scholar at LAAS-CNRS, Toulouse, (July to September 2015 and February 2016).

His research interests include control of UAVs and distributed optimization.



Alessandro Rucco received the Master's degree in automation engineering and the Ph.D. degree in information engineering from the Università del Salento, Lecce, Italy, in 2009 and 2014, respectively.

He was Visiting Student at the University of Paris Sud Supélec (March-April 2010) and at the University of Colorado Boulder (January-September 2012). He was Team Leader of the VI-RTUS team winning the International Student Competition Virtual Formula 2012. From August 2014 to September 2016, he was a Post-Doctoral Researcher at the

Department of Electrical and Computer Engineering, University of Porto, Portugal. From October 2016 to July 2018, he was a Post-Doctoral Researcher at the Department of Engineering, Università del Salento.

His research interests include distributed optimization, nonlinear optimal control, modeling, trajectory optimization, and high-performance maneuvering for car and aerial vehicles.



Giuseppe Notarstefano received the Laurea degree summa cum laude in electronics engineering from the Università di Pisa, Pisa, Italy, in 2003 and the Ph.D. degree in automation and operation research from the Università di Padova, Padua, Italy, in 2007.

He is a Professor with the Department of Electrical, Electronic, and Information Engineering G. Marconi, Alma Mater Studiorum Università di Bologna, Bologna, Italy. He was Associate Professor (from June 2016 to June 2018) and previously Assistant Professor, Ricercatore (from February 2007),

with the Università del Salento, Lecce, Italy. He has been Visiting Scholar at the University of Stuttgart, University of California Santa Barbara, Santa Barbara, CA, USA and University of Colorado Boulder, Boulder, CO, USA. His research interests include distributed optimization, cooperative control in complex networks, applied nonlinear optimal control, and trajectory optimization and maneuvering of aerial and car vehicles.

Dr. Notarstefano serves as an Associate Editor for *IEEE Transactions on Automatic Control*, *IEEE Transactions on Control Systems Technology*, and *IEEE Control Systems Letters*. He is also part of the Conference Editorial Board of IEEE Control Systems Society and EUCA. He is recipient of an ERC Starting Grant 2014.