Distributed Actuator Selection: Achieving Optimality via a Primal-Dual Algorithm

Licio Romao, Kostas Margellos, and Antonis Papachristodoulou

***Abstract*—This letter addresses the actuator selection problem, i.e., given an interconnection of asymptotically stable linear dynamical systems on a network and m possible actuators choose ν among them to achieve a certain objective. In general, this is a combinatorial optimization problem which is hard to solve; convex relaxations do not usually yield an optimal solution for the original problem. In this letter we focus on a particular instance of the actuator selection problem, namely the formulation with the trace of the controllability Gramian matrix as the optimization metric, and show that such a choice gives rise to an integer linear program (LP). Using properties of integral polyhedra, we show through a sequence of reformulations that the optimal solution of this problem can be determined by means of an LP without introducing any relaxation gap. This allows us to obtain the optimal solution using a primal-dual distributed algorithm, thus providing a scalable approach to the problem of actuator placement which has been up to now performed in a centralized manner enumerating all possible placement alternatives. We illustrate the main features of our approach by means of a case study involving a simplified model of the European power grid.**

***Index Terms*—Actuator selection problems, large-scale systems, optimization, distributed optimisation.**

I. Introduction

In THE past few years, an active stream of research within the control community has been understanding and regulating complex networks with applications to health care, neuroscience, and social networks [1]. There are various challenges in these attempts, however, achieving an implementable solution in the aforementioned areas requires scalability and reliability.

One important problem arising from the study of large-scale complex networks is that of actuator and sensor selection/placement. The former aims to choose ν from m potential actuator positions to minimise some objective function, e.g., $\mathcal{H}_2$ norm or controllability-related metrics, while the latter addresses a similar situation by deciding on ν sensors to minimise metrics related to estimation error. These problems have been extensively studied in the literature, ranging from applications in power systems [2] to the satellite assignment problem [3] and wireless networks [4].

Being combinatorial problems, actuator and sensor placement do not have, in general, efficient algorithms to determine their corresponding optimal solutions. The straightforward, albeit very commonly used, approach of enumerating all possible selections, and choosing the best one according to a given metric becomes intractable for high values of m. Even listing all possible alternatives requires shared information that in several applications might be considered as private.

An alternative approach is instead of obtaining the optimal solution, to rely on algorithms that possess guaranteed sub-optimality bounds, for instance, the greedy algorithm applied to submodular functions [5]–[7]. Such algorithms have been successfully applied to the actuator and sensor selection problems [8]–[11]. Furthermore, one can also sub-optimally solve the actuator and sensor placement problem using convex relaxations [2], [12]–[14], however, without *a priori* guaranteed sub-optimality bounds.

Within this context, we aim to study the actuator selection problem using the trace of the controllability gramian as optimisation metric. Under this metric and asymptotic stability of the dynamics, we show that the actuator placement problem can be equivalently posed as an Integer Linear Program (ILP). Using properties of integral polyhedra, we show through a sequence of reformulations that the optimal solution of this problem can be determined by means of a Linear Program (LP) without introducing any relaxation gap. This allows us to exploit recent results in [15], and to determine the optimal solution by means of a primal-dual distributed algorithm, thus providing a scalable approach to the problem of actuator placement which has been up to now performed on a centralised manner enumerating all possible placement alternatives.

Deviating from recent attempts in the literature, we recognize the combinatorial nature of the problem but do not rely on submodularity properties of set functions, as in [9], [10], and [16]. Our standpoint is closer to [12] since we use convex relaxations to study the problem. However, [12] focuses on different optimisation metrics and provides a sub-optimal solution to the problem. In contrast, this letter adopts a particular, in some sense simpler, optimisation metric, but

Manuscript received March 6, 2018; revised May 15, 2018; accepted June 7, 2018. Date of publication June 21, 2018; date of current version July 6, 2018. The work of L. Romao was supported by the Coordination for the Improvement of Higher Education Personnel (CAPES), Brazil, under the Grant 88881.128854/2016-01. The work of K. Margellos was supported by EPSRC U.K. under Grant EP/P03277X/1. The work of A. Papachristodoulou was supported by EPSRC U.K. under Grant EP/M002454/1. Recommended by Senior Editor G. Yin. *(Corresponding author: Licio Romao.)*

The authors are with the Department of Engineering Science, University of Oxford, Oxford OX1 3PJ, U.K. (e-mail: licio.romao@eng.ox.ac.uk; kostas.margellos@eng.ox.ac.uk; antonis@eng.ox.ac.uk).

Digital Object Identifier 10.1109/LCSYS.2018.2849611

obtains stronger results, showing that the optimal solution of the actuator selection problem can be obtained by means of a linear program. The modularity of the metric used opens the road for a distributed algorithmic implementation which is of particular interest in large-scale complex networks.

The remainder of this letter is organized as follows. Section II states the actuator selection problem under study, presents an equivalent ILP formulation, and provides an interpretation about the optimisation metric used in this letter. In Section III we introduce some background notions based on the properties of integral polyhedra and show that the optimal solution of the ILP can be obtained by means of a LP. We also introduce two algorithms, one decentralised (based on dual decomposition) and one distributed, that generate optimal solutions for the problem under study. In Section IV we illustrate the efficacy of our approach by means of a case study involving a simplified model of the European power grid.

II. The Actuator Selection Problem

A. Problem Statement

We consider the actuator selection problem using the trace of the controllability gramian as the optimisation metric. To this end, let m denote the number of possible actuators and $S \subset \{1, \dots, m\}$, and consider

$$\dot{x}(t) = Ax(t) + B_S u_S(t), \tag{1}$$

where $A \in \mathbb{R}^{n\times n}$ is a matrix that represents the dynamics of the system and $B_S \in \mathbb{R}^{n\times n_s}$ the input matrix associated with the subset S. The state of this system is denoted by $x(t) \in \mathbb{R}^n$ and $u(t) \in \mathbb{R}^{n_s}$ is the input. The objective is to choose a subset S of cardinality ν to maximise the trace of the controllability gramian, i.e.,

$$\underset{S\subset\{1,\dots,m\},\ |S|=\nu}{\text{maximise}} \quad \mathbf{tr}(W_S), \tag{2}$$

where $AW_S + W_S A^T + B_S B_S^T = 0$.

We impose the following standing assumption on the matrix A that describes the system's dynamics.

Assumption 1: Matrix A is Hurwitz, i.e., its eigenvalues have negative real part.

Note that, similar to [12], we can reformulate problem (2) as the Boolean-convex problem

$$\begin{aligned} \underset{W,z}{\text{maximise}} \quad & \mathbf{tr}(W) \\ \text{subject to} \quad & AW + WA^T + \sum_{i=1}^{m} z_i B_i B_i^T = 0 \\ & \mathbf{1}^T z = \nu, \quad z_i \in \{0,1\},\ \forall\, i = 1, \dots, m, \end{aligned} \tag{3}$$

by associating decision variables z_i, $i = 1, \dots, m$, to each actuator, where B_i is of appropriate dimension. The symbol $\mathbf{1}$ stands for the vector of ones in $\mathbb{R}^m$. Note that we allow B_i to be a matrix, not necessarily a vector, and do not impose any structure on B_i, $i = 1, \dots, m$. In the particular case where each actuator is connected to only one state we have $B_i = e_i$, for $i = 1, \dots, m$, where e_i is the standard unit vector.

Problem (3) can be simplified if we solve the Lyapunov equation. We can thus obtain the equivalent ILP formulation:

$$\begin{aligned} \underset{z\in\{0,1\}^m}{\text{minimise}} \quad & \sum_{i=1}^{m} c_i z_i \\ \text{subject to} \quad & \mathbf{1}^T z = \nu \end{aligned} \tag{$\mathcal{P}_{ILP}$}$$

where $c_i = -\mathbf{tr}(W_i)$, $AW_i + W_i A^T + B_i B_i^T = 0$. Note that equivalence follows from uniqueness of the solution to the Lyapunov equation, which is guaranteed by Assumption 1, while c_i, $i = 1, \dots, m$, can be computed in parallel for each actuator position.

We could formulate problem ($\mathcal{P}_{ILP}$) as one of minimising a modular[1] function over a cardinality constraint in the set of feasible solutions, which would reduce the search for the optimal solution of ($\mathcal{P}_{ILP}$) to a simple sorting problem. The study of combinatorial problems using modular and submodular functions has been the main subject in the recent literature [7]–[10], especially because of guaranteed sub-optimal bounds of the greedy algorithm [5], [6].

Even though in principle the solution to ($\mathcal{P}_{ILP}$) can be easily computed, this is not the case when m is large. Besides, privacy issues may also be a concern since some actuators may not be willing to share private data, e.g., the cost vector c_i. In this case, one could search for the optimal solution of the problem using distributed optimisation [17].

B. Optimisation Metric Interpretation

We provide an interpretation of our choice for the optimisation metric. To this end, note that maximising the trace of the controllability gramian results in maximising the sum of its eigenvalues, which can be thought of as a proxy for average controllability [10], [18]. At the same time, considering full state access, $\mathbf{tr}(W_S)$ coincides with the H_2 norm of (1). Therefore, by solving (2) we are maximising the energy of the impulse response of all possible actuator placement alternatives under the hypothesis of full-state measurement. This conclusion complies with the physical interpretation presented in [19] (see [10] and [20] for further details).

III. Distributed Implementation

A. ILP Background

Let us consider the following ILP

$$\begin{aligned} \underset{z\in\mathbb{Z}_+^n}{\text{minimise}} \quad & \sum_{i=1}^{m} c_i^T z_i \\ \text{subject to} \quad & \sum_{i=1}^{m} H_i z_i \le g, \end{aligned} \tag{4}$$

where $z = [z_1^T \dots z_m^T]$, with $z_i \in \mathbb{Z}_+^{n_i}$, $\sum_{i=1}^m n_i = n$, is the decision variable, $g \in \mathbb{Z}^p$ is the resource vector, and $\sum_{i=1}^m H_i z_i \le g$

[1]For $M \subset \mathbb{R}$, submodular function is a set function $f : 2^M \to \mathbb{R}$ with the property that $f(S \cup \{i\}) - f(S) \ge f(S \cup \{i,j\}) - f(S \cup \{j\}), \forall S \subset M \setminus \{i,j\}$. Intuitively, submodular functions have a diminishing return property, that is, the contribution of adding an element i deteriorates when the number of elements in S increases. Modular functions satisfy the previous inequality with equality.

is the coupling constraint, where $H_i \in \mathbb{Z}^{p\times n_i}$ – observe that $n_i = 1$ for all $i = 1,\dots,m$ in formulation $(\mathcal{P}_{ILP})$, which implies $n = m$. Instances of problem (4) include but are not limited to the knapsack and set covering problems [21]. Note that for the results of this subsection we allow z_i, $i = 1,\dots,m$, to be positive integer-valued but not necessarily binary as in $(\mathcal{P}_{ILP})$.

A non-negative integer vector z is a feasible solution of (4) if it satisfies the coupling constraint. The set of all feasible solutions is called the feasible set, and if non-empty, (4) is said to be feasible. We define the set of optimal solutions as the subset of the feasible set such that the value of the objective function is less than or equal to the value of any other vector in the feasible set.

In general, solving an ILP problem is hard because of the difficulty in characterizing the convex hull of the feasible set in terms of polyhedral inequalities [21]. In this direction, [22], [23] provide upper bounds on the difference between the optimal solution of (4) and its convex relaxation by tightening the resource vector g. However, in some special cases, we can produce a convex relaxation that is exact, i.e., its optimal solution produces an optimal solution for (4). For instance, the celebrated Birkhoff-von Neumann theorem [24] states that the extreme points of the set of doubly stochastic matrices are permutation matrices. Using this theorem one can solve the assignment problem, where we have m objects to assign to m people and aim to find the allocation with minimum cost. Note that the optimal solution for this problem is a permutation matrix, and it is often known as allocation in *merit order*. The Birkhoff-von Neumann theorem provides a way to produce an optimal solution by minimising over the set of doubly stochastic matrices instead of using the integer formulation with permutation matrices.

Towards this direction, integral polyhedra possess important properties that allow solving an ILP up to optimality by means of its convex relaxation.

Definition 1: A polyhedral set P in $\mathbb{R}^m$ is integral when all the extreme points of P have integer components.

The intuition behind Definition 1 is the result that a linear function that is not unbounded over P attains its minimum at some vertex of P [24, Ch. 2, Proposition 2.4.2]. We are implicitly assuming that P has at least one vertex, which is the case if and only if P does not contain a line [24, Ch. 2, Proposition 2.1.2]. A related definition in this context is the notion of total unimodularity.

Definition 2: Matrix $H \in \mathbb{Z}^{p\times m}$ is totally unimodular if the determinant of each submatrix is either 0, 1, or -1.

The following result, whose proof is given in [21, Ch. 3, Corollary 3.1], relates Definitions 1 and 2.

Lemma 1: The polyhedron $P = \{x \in \mathbb{R}^m_+ | Hx \le g, 0 \le x \le u\}$ is integral if and only if H is totally unimodular.

Note that, similar to the Birkhoff-von Neumann theorem, Lemma 1 provides mechanisms to exactly solve an ILP through its convex relaxation. Indeed, as an immediate consequence of Lemma 1, we can argue that whenever matrix $H = \begin{bmatrix} H_1 & \dots & H_m \end{bmatrix}$ is totally unimodular the optimal solution of problem (4) is equal to the optimal solution of its convex relaxation (i.e., considering $z \in \mathbb{R}^m_+$).

B. Convex Reformulation

Building on the results of Section III-A, we now focus on the actuator selection formulation given by $(\mathcal{P}_{ILP})$ and consider its convex relaxation

$$\begin{array}{ll} \underset{0\le z_i\le 1,\ \forall i=1,\dots,m}{\text{minimise}} & \sum_{i=1}^{m} c_i z_i \\ \text{subject to} & \mathbf{1}^T z = \nu. \end{array} \qquad (\mathcal{P}_{LP})$$

Proposition 1: The feasible set of $(\mathcal{P}_{ILP})$ coincides with extreme points of the polyhedron

$$P = \left\{ z \in \mathbb{R}^m | \mathbf{1}^T z = \nu, \quad 0 \le z_i \le 1, \quad \forall i = 1,\dots,m \right\},$$

i.e., the set of feasible solutions of $(\mathcal{P}_{LP})$, for all $\nu \in \mathbb{Z}_+$.

Proof: First, note that P does not contain a line since $0 \le z_i \le 1$ for all $i = 1,\dots,m$, hence it has at least one extreme point. Then, define H as previously with

$$H_i = \begin{bmatrix} 1 \\ -1 \end{bmatrix}, \quad \forall i = 1,\dots,m, \text{ and } g = \begin{bmatrix} \nu \\ -\nu \end{bmatrix}. \qquad (5)$$

Now, let $\hat{H}$ be a square submatrix of H. We have two possible cases: $\hat{H}$ has dimension 1 or 2. In the former case, note that either $\hat{H} = 1$ or $\hat{H} = -1$, which trivially satisfies $\det(\hat{H}) = \pm 1$, where $\det(A)$ denotes the determinant of A. In the latter case, observe that every 2×2 submatrix has determinant zero. Therefore, by Definition 1, H is totally unimodular. The result follows directly from Lemma 1, thus concluding the proof. ∎

C. Primal-Dual Algorithms

Proposition 1 shows that an optimal solution of $(\mathcal{P}_{ILP})$ can be recovered by means of $(\mathcal{P}_{LP})$. However, the latter is a linear program. In this section we exploit this fact, and provide a decentralised and a distributed algorithm to obtain its optimal solution. To this end, we equivalently represent the constraints of $(\mathcal{P}_{LP})$ as $Hz \le g$, where H and g are given by (5). The corresponding dual problem is given by

$$\underset{\lambda\ge 0}{\text{maximise}} \quad -\lambda^T g + \sum_{i=1}^{m} \min_{0\le z_i\le 1} (c_i + \lambda^T H_i) z_i \qquad (\mathcal{D}_{LP})$$

The minimisation step required to evaluate the dual function can be performed in parallel for each $i = 1,\dots,m$, making this formulation amenable to decomposition algorithms. Second, we have zero duality gap between problems $(\mathcal{D}_{LP})$ and $(\mathcal{P}_{LP})$ (by strong duality arguments in linear programming), which in turn provides an optimal solution for its integer formulation given in $(\mathcal{P}_{ILP})$.

One way to solve problem $(\mathcal{D}_{LP})$ is to use dual ascent methods, which are encoded by

$$\begin{aligned} z_i^{k+1} &\in \underset{0\le z\le 1}{\operatorname{argmin}} (c_i + (\lambda^k)^T H_i) z, \quad i = 1,\dots,m \\ \lambda^{k+1} &\in \left[\lambda^k + \alpha^k \left(\sum_{i=1}^{m} H_i z_i^{k+1} - g \right) \right]_+, \end{aligned} \qquad (6)$$

where $[\cdot]_+$ represents the projection of its argument on the positive orthant. Typical choices for the time-varying step-size are

Algorithm 1 Distributed Algorithm of [15]

Require: : $g, H_i, c_i, \quad \forall i = 1, \ldots, m$

$z_i^0 \in [0,1], \quad \forall i = 1, \ldots, m$

$\lambda_i^0 = 0, \quad \forall i = 1, \ldots, m$

$k = 0$

while convergence is not achieved **do**

  for $i = 1$ to m **do**

    $\ell_i^k = \sum_{j=1}^m a_{ij}^k \lambda_j^k$

    $z_i^{k+1} \in \operatorname{argmin}_{0\le z\le 1}(c_i + (\ell_i^k)^T H_i) z$

    $\lambda_i^{k+1} \in [\ell_i^k + \alpha^k (H_i z_i^{k+1} - \frac{g}{m})]_+$

    $\hat{z}_i^{k+1} = \hat{z}_i^k + \frac{\alpha^k}{\sum_{r=1}^k \alpha^r}(z_i^k - \hat{z}_i^k)$

  end for

  $k \leftarrow k+1$

end while

$\alpha^k = \beta/(k+1)$ or $\alpha^k = \beta/\sqrt{k+1}$, for some $\beta > 0$. As shown in [25], with either choice for α^k, the sequence generated in (6) converges to the set of optimal solutions of $(\mathcal{D}_{LP})$.

Notwithstanding the progress made from the centralised problem $(\mathcal{P}_{LP})$ to the decentralised Algorithm (6), this approach still has some drawbacks. At each iteration, the dual variable λ needs to be broadcast in order to perform the next primal update, which must be sent to the central processor that performs the dual update.

To alleviate the centralised dual update step we introduce a time-varying communication network with edge weights a_{ij}^k, where $a_{ij}^k = 0$ implies that node j does not share information with node i at iteration k.

Assumption 2 (Weight Coefficients): For each $k \ge 0$, $a_{ij}^k \in [0, 1)$ and there exists $\eta \in (0, 1)$ such that $a_{ii}^k \ge \eta$ and $a_{ij}^k > 0$ implies $a_{ij}^k \ge \eta$. Moreover, we have $\sum_{i=1}^m a_{ij}^k = 1, \ \forall i = 1, \ldots, m$ and $\sum_{j=1}^m a_{ij}^k = 1, \ \forall j = 1, \ldots, m$.

In other words, Assumption 2 requires the weight matrix to be doubly stochastic with the property that when agents i and j communicate then the exchanged information should be assigned a non-zero weight, i.e., $a_{ij}^k > 0$ implies $a_{ij}^k \ge \eta$. Besides, we need to impose certain connectivity requirements. At each iteration, we define a directed graph, (V, E_k), where $V = \{1, \ldots, m\}$ represents the nodes and $E_k = \{(j, i) : a_{ij}^k > 0\}$ the edges, to capture the induced structure of the communication network. We also define the matrix $E_\infty = \{(j, i) : (j, i) \in E_k$ for infinitely many $k\}$. We then impose the following assumption.

Assumption 3 (Connectivity and Communication): Graph (V, E_∞) is strongly connected and there exists $T \ge 1$ such that for every $(j, i) \in E_\infty$ agent i receives information from a neighbouring agent j at least once every consecutive T iterations.

We are now in a position to present Algorithm 1, which was proposed in [15], and in contrast to (6) offers a distributed implementation. We will henceforth refer to each node i, $i = 1, \ldots, m$, as agent, which interacts with other agents according to the aforementioned communication protocol. By Algorithm 1, each agent calculates a weighted sum of the dual variables that were calculated by neighbouring agents according to the underlying communication structure at the previous iteration (Step 6) to update its own dual variable (Step 8), eliminating the need for a central agent to perform the dual update. Additionally, observe that variable α^k is a time-varying step size, with the following property.

Assumption 4 (Time-Varying Step Size): The sequence $\{\alpha^k\}_{k\ge 0}$ is non-increasing, positive, $\sum_{k=1}^\infty \alpha^k = \infty$, and $\sum_{k=1}^\infty (\alpha^k)^2 < \infty$.

A common choice for the time-varying step is $\alpha^k = \frac{\beta}{k+1}$, for some $\beta > 0$. Assumptions 2, 3 and 4 are essential for the convergence proof (see [15] for details).

Note that in both Algorithms 1 and (6) the primal update is performed by the following "if then else" clause

$$z_i(\zeta) = \operatorname*{argmin}_{0\le z\le 1}(c_i + \zeta^T H_i) z = \begin{cases} 1, & c_i + \zeta^T H_i \le 0, \\ 0, & \text{otherwise}, \end{cases}$$

where $\zeta = \lambda^k$ in (6) and $\zeta = \ell_i^k$ in Algorithm 1. As a consequence, Step 7 in Algorithm 1 and the primal update in (6) is computationally inexpensive.

Under assumptions 2, 3 and 4, it is shown in [15] that $\lim_{k\to\infty} \operatorname{dist}(\hat{z}^k, Z^\star)$ is equal to zero, where $\hat{z}^k$ is the sequence generated by Algorithm 1, $Z^\star$ is the set of optimal solution of $(\mathcal{P}_{LP})$ and $\operatorname{dist}(x, X)$ is the distance between the point x and the set X.

Theorem 1: Consider the actuator selection formulation given by (2) and suppose Assumptions 2–4 hold. Then the sequence $\hat{z}^k$ generated by Algorithm 1 produces the optimal objective value for (2).

Proof: By the proof of [15, Th. 2] the sequence $\{\hat{z}^k\}_{k\ge 0}$ converges to the set of optimal solutions of $(\mathcal{P}_{LP})$, which, through Proposition 1, implies that the optimal objective function values of $(\mathcal{P}_{LP})$ and $(\mathcal{P}_{ILP})$ coincide. However, by the equivalence between $(\mathcal{P}_{ILP})$ and (2) (see Section II) the result follows. ■

The proof of Theorem 1 relies on several equivalences. These can be summarised as

$$(2) \overset{\star}{\iff} (\mathcal{P}_{ILP}) \overset{Prop.\ 1}{\iff} (\mathcal{P}_{LP}) \overset{\#}{\iff} (\mathcal{D}_{LP}).$$

The left-hand side equivalence ($\star$) stands for the formulation of the actuator selection problem as an ILP through the steps presented in Section II. Some ideas used in this process appear in [12] and [16], but the last step where we convert into an ILP had not been explored in the literature. The right-hand side equivalence (#) is due to linear programming strong duality (see [15] and [17]). The intermediate, instrumental equivalence result, was established by means of Proposition 1.

Under uniqueness[2] of the solution of $(\mathcal{P}_{LP})$, we can strengthen the result of Theorem 1.

Corollary 1: Assume $(\mathcal{P}_{LP})$ admits a unique solution and let $\xi = \lim_{k\to\infty} \hat{z}^k$. We then have that the optimal solution $S^\star$ of (2) is given by $S^\star = \{i | \xi_i = 1\}$.

Proof: According to Theorem 1, the optimal values of $(\mathcal{P}_{LP})$ and (2) are the same. Since the solution of $(\mathcal{P}_{LP})$ is unique, we know that $\xi = \lim_{k\to\infty} \hat{z}^k$ is well-defined and, in addition, it is an extreme point of the set of feasible solutions of $(\mathcal{P}_{LP})$ (by Proposition 1). Hence, it is a feasible solution for (2).

[2]Even without uniqueness of the solution of $(\mathcal{P}_{LP})$ we can always produce an optimal solution for (2), however, this solution would differ according to the iteration index (corresponding to practical convergence) where Algorithm 1 or iteration (9) is terminated.

Due to Theorem 1 it is also optimal since ξ will achieve the optimal value of $(\mathcal{P}_{LP})$. By the equivalence of $(\mathcal{P}_{LP})$ and (2), $S^\star$ constitutes the optimal solution of (2), thus concluding the proof. ■

Remark 1: One can prove that a similar result holds for the decentralised algorithm in (6), namely, at the limit, the sequence z^k generated by (6) converges in terms of value to the optimal value of (2). For the sake of brevity, we do not present the proof here, however, it follows from the proof of Theorem 1 by setting the weights in Assumption 2 to be iteration invariant, and all of them equal to $1/m$. Moreover, it is worth pointing out that our main result does not depend on the particular algorithmic choice, and other distributed or decentralised schemes could be employed as well.

IV. Numerical Example: European Power Grid

In this section, we illustrate the proposed method for the actuator selection problem using a case study involving a simplified model of the European power grid to decide the location of HVDC links in the network.

In particular, we revisit the HVDC allocation problem studied in [10]. In general, HVDC links are employed to enhance transient response of the power system by influencing active and reactive power injections to damp frequency oscillations and prevent rotor angle instability [10], [26]. The model consists of a linear system that represents the European grid, which is composed of 74 buses connected to a generator and a constant impedance load. The linear model is obtained after linearising the swing equations about nominal operating points for each possible HVDC link placement. As in [10], the purpose of this example is to assess the efficacy of the proposed algorithm in a realistic setting, while from an application point of view further investigation is needed, as in HVDC placement decisions in the European power system controllability is just one of several objectives, e.g., economic. More details about the model can be found in [2] and [26], and references therein.

The linearised model has 148 states, since each bus consists of one generator and each generator has two state variables corresponding to rotor angle and frequency dynamics. Following [10] we suppose that any generator can be possibly connected to any other generator in the grid, which yields 2701 possible connections, from which we want to select the best 10 placements according to the controllability trace optimisation metric. Simple calculations show that this configuration gives us a total of approximately 5.6×10^{27} possibilities, which is far beyond a sorting algorithm enumerating all alternatives can handle.

To implement Algorithm 1 all agents receive matrix A, and each agent its own matrix B_i, $i = 1, \ldots, 2701$, from the network. Upon this, agents can locally compute their own c_i by solving the corresponding Lyapunov equation.

The convergence rate of Algorithm 1 can be improved by introducing a sequence $\tilde{z}_i^k$ defined as

$$\tilde{z}_i^{k+1} = \begin{cases} \hat{z}_i^{k+1}, & k < L, \\ \dfrac{\sum_{r=L}^{k} \alpha^r z_i^{k+1}}{\sum_{r=L}^{k} \alpha^r}, & k \geq L, \end{cases}$$

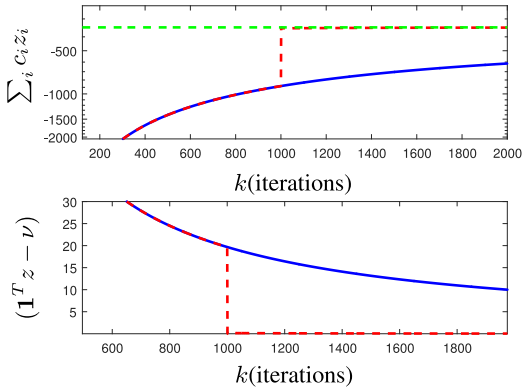

Fig. 1. Evolution of the objective function (top graph) and constraint violation (bottom graph) for the sequences $\hat{z}^k$ (solid blue line) and $\tilde{z}^k$ (dashed red line) of Algorithm 1 for the European power model with $\alpha^k = \frac{10}{(k+1)}$. In the top graph, the dashed green line corresponds to the solution of the centralised problem counterpart.

where L is defined to be a fraction of the total number of iterations. Similar arguments used to prove convergence of sequence $\{\hat{z}\}_{k\geq 0}$ of Algorithm 1 to the optimal solution of $(\mathcal{P}_{LP})$ applies to $\{\tilde{z}\}_{k\geq 0}$ [15, Th. 2]. The latter sequence alleviates the influence of bad estimates of the dual variable in earlier steps of Algorithm 1.

In our simulations, we initialise the primal and dual variables to zero, as suggested in steps 1 and 2 of Algorithm 1. Furthermore, we define the iteration-varying step size as $\alpha^k = \beta/(k+1)$, with $\beta = 10$, and run 2000 iterations. We also create a time-varying communication structure to satisfy Assumptions 2 and 3 by alternating between two strongly connected graphs, i.e., a_{ij}^{2r} and a_{ij}^{2r+1} are constant for all $r \in \mathbb{Z}_+$. To guarantee conditions of Assumption 2 for each graph, we generate a doubly stochastic matrix by forming a convex combination of 100 randomly generated permutation matrices, making sure we include the identity matrix to ensure $a_{ii}^k \geq \eta$ with $\eta \in (0, 1)$. If the generated matrix satisfies $a_{ij}^k \geq \eta$ whenever $a_{ij}^k > 0$ we return the matrix; in the negative case, we repeat the process until these assumptions are achieved. Besides, if this algorithm terminates, Assumptions 2 and 3 are satisfied with $T = 2$.

Figure 1 shows the evolution of the objective function (top graph) and the constraint violation for the sequences $\hat{z}^k$ (solid blue line) and $\tilde{z}^k$ (dashed red line). In the top graph, the optimal value of the centralised problem counterpart is shown by means of a dashed green line. Note that, at the beginning, the algorithm exhibits superior performance as far as the optimal value is concerned, however, the generated primal iterates are infeasible. As the algorithm progresses, we degrade the value of the objective to achieve primal feasibility. Additionally, as we can see, the sequence $\tilde{z}^k$ has a better convergence rate than sequence $\hat{z}^k$ (we use $L = 1000$ in our simulation).

For completeness, we also apply the decentralised algorithm whose main steps are encoded by (9). As a stopping criterion for Algorithm 1 we use primal feasibility. Figure 2 illustrates the results with the choice of $\alpha^k = 0.1/\sqrt{(k+1)}$ for the time-vanishing step size. Observe that the optimal solution is achieved when we find a primal feasible solution, which occurs around 3640 iterations (see zoomed areas).

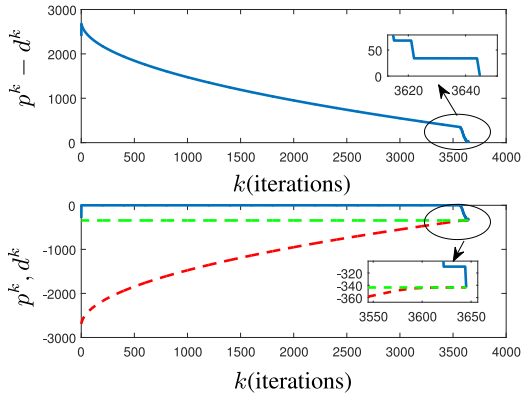

Fig. 2. Duality gap (top) and values of the primal and dual objective values (bottom) for the dual algorithm given in (9). In the bottom graph, the primal objective value at each iteration is denoted by the solid blue line; the dual objective by the dashed red line; and the centralised optimal by the dashed green line. A primal feasible solution is found after approximately 3640 iterations.

Both Algorithm 1 and (6) converge to the optimal solution of the problem, which in this case admits a unique solution (see Corollary 1), the former being more adequate for large scale networks because it does not need the dual variable to be updated by a central processor. It should be mentioned that the parameter β of the time-varying step-size, albeit not interfering in the theoretical convergence, is decisive for the rate of convergence.

V. Conclusion

Using properties of integral polyhedra, we showed that the optimal solution of the actuator selection problem can be obtained by means of a linear program when the trace of the controllability gramian is employed as the optimisation metric. To solve the resulting LP, we provided two primal-dual algorithms: a decentralised one based on dual decomposition; and a distributed one that does not require a common processor for dual updates. The efficacy of our methods was investigated numerically on a European power grid case study.

Interesting extensions include dealing with constraints in the control input and increasing the privacy level in the communication between agents. Besides, it is worth noting that our approach relies on the fact that the optimal solution using the trace of the gramian as optimisation metric is on the vertices of a polyhedron. Other, more general, metrics can be handled by means of convex relaxations [2], [12]; or by greedy algorithms if the objective is submodular [10].

Acknowledgment

The authors would like to thank Tyler Summers for providing the model of the European power grid used in Section IV.

References

[1] F. Lamnabhi-Lagarrigue *et al.*, "Systems & control for the future of humanity, research agenda: Current and future roles, impact and grand challenges," *Annu. Rev. Control*, vol. 43, pp. 1–64, Jan. 2017.

[2] A. Fuchs and M. Morari, "Placement of HVDC links for power grid stabilization during transients," in *Proc. IEEE Grenoble Conf. PowerTech (POWERTECH)*, Grenoble, France, Nov. 2013, pp. 1–6.

[3] G. Qu, D. Brown, and N. Li, "Distributed greedy algorithm for satellite assignment problem with submodular utility function," *IFAC Papers Online*, vol. 48, no. 22, pp. 258–263, Sep. 2015.

[4] V. Gupta, T. H. Chung, B. Hassibi, and R. M. Murray, "On a stochastic sensor selection algorithm with applications in sensor scheduling and sensor coverage," *Automatica*, vol. 42, no. 2, pp. 251–260, 2006.

[5] L. Lovász, "Submodular functions and convexity," in *Mathematical Programming the State of the Art*, A. Bachem, B. Korte, and M. Grötschel, Eds. Berlin, Germany: Springer, 1983.

[6] G. L. Nemhauser and L. A. Wolsey, "Maximizing submodular set functions: Formulations and analysis of algorithms," *North Holland Math. Stud.*, vol. 59, pp. 279–301, Jan. 1981.

[7] A. Krause and D. Golovin, "Submodular function maximization," in *Tractability Practical Approaches to Hard Problems*, L. Bordeaux and Y. Hamadi, Eds. Cambridge, U.K.: Cambridge Univ. Press, 2014, pp. 71–104.

[8] A. Krause, A. Singh, and C. Guestrin, "Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies," *J. Mach. Learn. Res.*, vol. 9, pp. 235–284, Feb. 2008.

[9] S. T. Jawaid and S. L. Smith, "Submodularity and greedy algorithms in sensor scheduling for linear dynamical systems," *Automatica*, vol. 61, pp. 282–288, Nov. 2015.

[10] T. H. Summers, F. L. Cortesi, and J. Lygeros, "On submodularity and controllability in complex dynamical networks," *IEEE Trans. Control Netw. Syst.*, vol. 3, no. 1, pp. 91–101, Mar. 2016.

[11] A. Clark, B. Alomair, L. Bushnell, and R. Poovendran, "Submodularity in input node selection for networked systems linear systems: Efficient algorithms for performance and controllability," *IEEE Control Syst.*, vol. 37, no. 6, pp. 52–74, Dec. 2017.

[12] S. Joshi and S. Boyd, "Sensor selection via convex optimization," *IEEE Trans. Signal Process.*, vol. 57, no. 2, pp. 451–462, Feb. 2009.

[13] B. Polyak, M. Khlebnikov, and P. Shcherbakov, "An LMI approach to structured sparse feedback design in linear control systems," in *Proc. Eur. Control Conf.*, Zürich, Switzerland, Jul. 2013, pp. 833–838.

[14] N. K. Dhingra, M. R. Jovanović, and Z.-Q. Luo, "An ADMM algorithm for optimal sensor and actuator selection," in *Proc. 53rd IEEE Conf. Decis. Control*, Los Angeles, CA, USA, Dec. 2014, pp. 4039–4044.

[15] A. Falsone, K. Margellos, S. Garatti, and M. Prandini, "Dual decomposition for multi-agent distributed optimization with coupling constraints," *Automatica*, vol. 84, pp. 149–158, Oct. 2017.

[16] T. Summers and I. Shames, "Convex relaxations and Gramian rank constraints for sensor and actuator selection in networks," in *Proc. IEEE Int. Symp. Intell. Control*, Buenos Aires, Argentina, Sep. 2016, pp. 341–346.

[17] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Belmont, MA, USA: Athena Sci., 1989.

[18] K. Fitch and N. E. Leonard, "Optimal leader selection for controllability and robustness in multi-agent networks," in *Proc. Eur. Control Conf.*, Aalborg, Denmark, Jun. 2016, pp. 1550–1555.

[19] S. Leleu, H. Abou-Kandil, and Y. Bonnassieux, "Piezoelectric actuators and sensors location for active control of flexible structures," in *Proc. 17th IEEE Instrum. Meas. Technol. Conf. (IMTC)*, Baltimore, MD, USA, 2000, pp. 818–823.

[20] F. Pasqualetti, S. Zampieri, and F. Bullo, "Controllability metrics, and algorithms for complex networks," *IEEE Trans. Control Netw. Syst.*, vol. 1, no. 1, pp. 40–52, Mar. 2014.

[21] D. Bertsimas and R. Weismantel, *Optimization Over Integers*. Belmont, MA, USA: Dyn. Ideas, 2005.

[22] 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, pp. 144–156, May 2016.

[23] A. Falsone, K. Margellos, and M. Prandini. (Jun. 2017). *A Decentralized Approach to Multi-Agent MILPs: Finite-Time Feasibility and Performance Guarantees*. [Online]. Available: http://arxiv.org/abs/1706.08788

[24] D. P. Bertsekas, *Convex Optimization Theory*. Belmont, MA, USA: Athena Sci., 2009.

[25] D. P. Bertsekas, *Nonlinear Programming*, 2nd ed. Belmont, MA, USA: Athena Sci., 1999.

[26] A. Fuchs, S. Mariéthoz, M. Larsson, and M. Morari, "Grid stabilization through VSC-HVDC using wide area measurements," in *Proc. IEEE PES Trondheim PowerTech (POWERTECH)*, Trondheim, Norway, Jun. 2011, pp. 1–6.