k-Dimensional Agreement in Multiagent Systems

Gianluca Bianchin, Miguel Vaquero, Jorge Cortés, and Emiliano Dall'Anese

***Abstract*— We study the problem of k-dimensional linear agreement, whereby a group of agents is interested in computing k independent weighted means of a global vector whose entries are known only by individual agents. This problem is motivated by applications in distributed computing and sensing, where agents seek to evaluate multiple independent functions at a common vector point by running a single distributed algorithm. We propose the use of linear network protocols for this task, and we show that linear dynamics can agree on quantities that are oblique projections of the global vector onto certain subspaces. Moreover, we provide algebraic necessary and sufficient conditions that characterize all agreement protocols that are consistent with a certain graph, we propose a design procedure for constructing such protocols, and we study what classes of graphs can achieve agreement on arbitrary weights. Overall, our results suggest that k-dimensional agreement requires the use of communication graphs with higher connectivity compared to standard consensus algorithms; more precisely, we relate the existence of Hamiltonian decompositions in a graph with the capability of that graph to sustain an agreement protocol. The applicability of the framework is illustrated via simulations for two problems in robotic formation and in distributed regression.**

I. Introduction

COORDINATION and consensus protocols are central to many network synchronization problems, including rendezvous, distributed convex optimization, and distributed computation and sensing. An extensive literature is available on consensus-based processes, whereby the states of the network nodes asymptotically compute a *common* value that is a weighted average of the initial estimates available locally at the agents (see, for example, the representative works in context [1]–[3]). On the other hand, in several applications it is instead of interest to compute, asymptotically, $k > 1$ *independent* weighted averages (here, "independent" means that the vectors of weights are linearly independent) of the initial estimates. Relevant examples of this problem include distributed computation [4], [5], where the weighting accounts for different desired outcomes, task allocation algorithms [6], where weights account for heterogeneous computational capabilities, distributed sensing [7], [8], where the weighting is proportional to the accuracy of each sensing device, and robotic formation [9], where one agent would like to achieve a certain configuration *relative to* another agent.

G. Bianchin and E. Dall'Anese are with the Department of Electrical, Computer, and Energy Engineering, University of Colorado Boulder. M. Vaquero is with the School of Science and Technology, IE University. J. Cortés is with the Department of Mechanical and Aerospace Engineering, University of California San Diego.

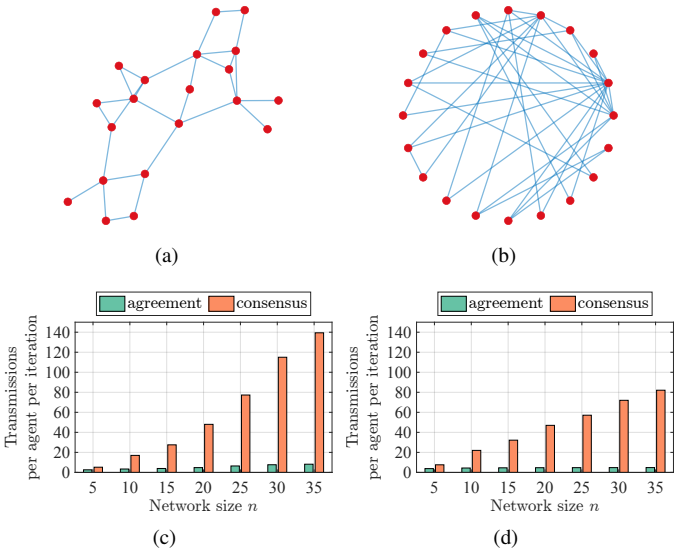

Fig. 1: Communication complexity of running k consensus algorithms in parallel vs one k-dimensional agreement algorithm (proposed in this paper) to compute $k = \lfloor \frac{n}{2} \rfloor$ weighted average means of a global quantity. (a) and (c) Erdős–Rényi network model. (b) and (d) Barabasi-Albert model. Bars denote the average number of transmissions per iteration per agent. See Example 3.4.

Mathematically, given a vector $x_0 \in \mathbb{R}^n$ of initial states or estimates – with each entry available locally at a single agent – and an arbitrary matrix $W \in \mathbb{R}^{n \times n}$ of rank $1 \leq k \leq n-1$ describing the desired weights, we say that the group reaches a k-dimensional agreement if, asymptotically, the agents' states converge to Wx_0. We model the communication between agents using a directed graph; our goal is to identify what classes of graphs are sufficiently rich to enable k-dimensional agreement and to determine, when possible, a distributed agreement protocol compatible with such a graph.

A natural approach to this problem consists in executing k classical consensus algorithms [2] in parallel, each designed to converge to one of the independent weighted means. However, the communication and computational complexities of such an approach do not scale with the network size; thus, our objective here is to achieve agreement by running a *single distributed algorithm*. A comparison of the communication load involved by k parallel consensus algorithms and the proposed k-dimensional agreement method is exemplified in Fig. 1 (simulation details are provided in Example 3.4).

Related work. The agreement protocol studied in this work

stems from the problem of distributed consensus. Because of their centrality, consensus algorithms have been extensively studied in the literature. A list of representative topics (necessarily incomplete) includes: sufficient and/or necessary conditions to reach consensus are provided in [2], [10]–[14], time delays are accounted in [12], consensus with arbitrary objective maps is studied in [15], the alternating direction method of multipliers (ADMM) is used in [16], [17] with linear ADMM-based protocols studied in [18], convergence rate is considered in [19], [20], and robustness in [21], [22]. Differently from constrained consensus problems [23], [24], where the agents' states must satisfy agent-dependent constraints during transients and the asymptotic value lies in the intersection of the constraint sets, in our setting the values are instead constrained at convergence, and thus state of the agents may not converge to identical values. In Pareto optimal distributed optimization [25], the group of agents cooperatively seeks to determine the minimizer of a cost function that depends on agent-dependent decision variables. Clustering-based consensus [26]–[28] is a closely-related problem where the states of agents in the same cluster of the graph converge to identical values, while inter-cluster states can differ. Differently from this setting, which is obtained by using weakly-connected communication graphs to separate the state of different communities, in this work we focus on cases where the asymptotic state of each agent depends on every other agent in the network. To the best of our knowledge, agreement problems exhibiting this dependence where the agents' states do not converge to identical values have not been studied. A relevant exception is the problem of scaled consensus considered in [29], where agents agree on subspaces of dimension $k = 1$. In this paper, we tackle the more general problem $k \geq 1$; as shown shortly below, the extension to $k > 1$ is non-trivial since standard assumptions made for consensus are not sufficient to guarantee agreement (see Example 4.5).

Contributions. The contribution of this work is fourfold. First, we formulate the problem of k-dimensional agreement and we propose the use of continuous-time linear network protocols to achieve this objective. We then show that by using linear protocols the set of agents can compute, asymptotically, points that are oblique projections of the vector of initial estimates. Conversely, we also show that, given any desired oblique projection of the initial estimates, it is always possible to design a corresponding agreement protocol provided that the underlying communication graph is sufficiently connected. In this case, we provide an algorithm to design agreement protocols. Second, for sparse communication graphs, we provide an algebraic characterization of all protocols (consistent with a certain communication graph) that guarantee agreement. Further, we show how such conditions can be used to design efficient numerical algorithms that yield fast agreement. Third, by using the stated characterization, we provide graph-theoretic necessary and sufficient conditions to check whether a graph can sustain agreement dynamics. Our necessary conditions illustrate that widely-adopted graphs, such as the line and circulant topologies, can admit agreement protocols on subspaces of dimension at most $k = 1$. Our sufficient conditions show that graphs that admit independent Hamiltonian decompositions can always sustain agreement protocols on arbitrary weights. Fourth, we show that agreement algorithms can be adapted to account for cases where the local estimates are time-varying and, in this case, we prove convergence of these algorithms in the form of an input-to-state stability-type bound. Finally, we illustrate the applicability of the framework on regression and robotic coordination problems through a set of simulations.

Organization. Section II introduces basic concepts and Section III formalizes the problem. Section IV studies agreement protocols over complete graphs, Section V illustrates our algebraic characterization of agreement protocols and illustrates numerical methods to compute agreement algorithms, Section VI provides graph-theoretic conditions for agreement. Section VII extends the approach to tracking problems and Section VIII illustrates the techniques via numerical simulations. Conclusions and directions are discussed in Section IX.

II. Preliminaries

We first introduce some basic notions used in the paper.

Notation. We denote by $\mathbb{N}_{>0} = \{1, 2, \dots\}$ the set of positive natural numbers, by $\mathbb{C}$ the set of complex numbers, and by $\mathbb{R}$ the set of real numbers. Given $x \in \mathbb{C}$, $\Re(x)$ and $\Im(x)$ denote its real and imaginary parts, respectively. Given vectors $x \in \mathbb{R}^n$ and $u \in \mathbb{R}^m$, we let $(x, u) \in \mathbb{R}^{n+m}$ denote their concatenation. We denote by $\mathbb{1}_n \in \mathbb{R}^n$ the vector of all ones, by $I_n \in \mathbb{R}^{n\times n}$ the identity matrix, by $\mathbb{0}_{n,m} \in \mathbb{R}^{n\times m}$ the matrix of all zeros (subscripts are dropped when dimensions are clear from the context). Given $A \in \mathbb{R}^{n\times n}$, we denote its spectrum by $\sigma(A) = \{\lambda \in \mathbb{C} \ : \ \det(\lambda I - A) = 0\}$, and by $\lambda_{\max}(A) = \max\{\Re(\lambda) \ : \ \lambda \in \sigma(A)\}$ its spectral abscissa; also, we use the notation $A = [a_{ij}]$, where a_{ij} is the element in row i and column j of A. Given $A \in \mathbb{R}^{n\times m}$, $\mathrm{Im}(A)$ and $\ker(A)$ denote its image and null space, respectively. Given a polynomial with real coefficients $p(\lambda) = \lambda^n + p_1\lambda^{n-1} + \cdots + p_n$, $p(\lambda)$ is *stable* if all its roots have negative real part.

Graph-theoretic notions. A directed graph (abbreviated digraph), denoted by $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, consists of a set of nodes $\mathcal{V} = \{1, \dots, n\}$ and a set of directed edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$. An element $(i, j) \in \mathcal{E}$ denotes a directed edge from node j to i. We will often use the notion of weighted digraph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, A)$, where $A \in \mathbb{R}^{n\times n}$ is the graph's adjacency matrix; $A = [a_{ij}]$ satisfies $a_{ij} \neq 0$ only if $(i, j) \in \mathcal{E}$ and $a_{ij} = 0$ if $(i, j) \notin \mathcal{E}$. More generally, for fixed $\mathcal{G}$ and $A \in \mathbb{R}^{n\times n}$, we say that A is *consistent* with $\mathcal{G}$ if A is has the zero-nonzero pattern of an adjacency matrix for $\mathcal{G}$. The set of (in)*neighbors* of $i \in \mathcal{V}$ is $\mathcal{N}_i = \{j \in \mathcal{V} \setminus \{i\} \ : \ (i, j) \in \mathcal{E}\}$. A graph is *complete* if there exists an edge connecting every pair of nodes, and is *sparse* otherwise. A *path* in $\mathcal{G}$ is a sequence of edges $(e_1, e_2, \dots)$, $e_j \in \mathcal{E}$ for all $j = 1, 2, \dots$, such that the initial node of each edge is the final node of the preceding edge. Notice that a path may contain repeated edges and, also, going along the path one may reach repeated nodes. The *length* of a path is the number of edges contained in the sequence $(e_1, e_2, \dots)$. A graph is *strongly connected* if, for any $i, j \in \mathcal{V}$, there is a directed path from i to j. A *closed path* is a path whose initial and

final vertices coincide. A closed path is a *cycle* if, going along the path, one reaches no node, other than the initial-final node, more than once. A cycle of length equal to one is a *self cycle*. We say that $\mathcal{G}_1, \dots, \mathcal{G}_m \subseteq \mathcal{G}$ is a *decomposition* of $\mathcal{G}$ if the $\mathcal{G}_i$-s are pairwise disjoint and the union of the node sets of the $\mathcal{G}_i$-s is the node set of $\mathcal{G}$. A *Hamiltonian cycle* is a cycle that visits every node of $\mathcal{G}$ exactly once [30]. A *Hamiltonian decomposition* is a decomposition of $\mathcal{G}$ such that each subgraph $\mathcal{G}_i$ is a Hamiltonian cycle. (Notice that not all graphs admit a Hamiltonian decomposition; conversely, certain graphs admit multiple Hamiltonian decompositions.) A *Hamiltonian ℓ-decomposition* is a Hamiltonian decomposition of some subgraph of $\mathcal{G}$ with ℓ nodes, see Fig. 2 for illustration.

Structural analysis and sparse matrix spaces. We rely on the structural approach to systems theory proposed in [31], [32]. We are concerned with linear subspaces obtained by forcing certain entries of the matrices in $\mathbb{R}^{n\times n}$ to be zero. More precisely, given a digraph $\mathcal{G}$, we define $\mathcal{A}_\mathcal{G}$ to be the vector space in $\mathbb{R}^{n\times n}$ that contains all matrices that are consistent with $\mathcal{G}$, and whose nonzero entries are independent parameters. The vector space $\mathcal{A}_\mathcal{G}$ can be represented by a $n \times n$ *structure matrix* $\mathbf{A}_\mathcal{G}$, whose entries are either algebraically independent parameters in $\mathbb{R}$ (denoted by $\divideontimes$) or fixed zeros (denoted by 0). We define a vector of parameters $a = (\{a_{ij}\}_{(i,j)\in\mathcal{E}})$ such that $\mathbf{A}_\mathcal{G}(a)$ defines a numerical realization of the structured matrix $\mathbf{A}_\mathcal{G}$, namely, $\mathbf{A}_\mathcal{G}(a) \in \mathcal{A}_\mathcal{G}$. For instance, for the graph in Fig. 2, we have

$$\mathbf{A}_\mathcal{G} = \begin{bmatrix} \divideontimes & 0 & \divideontimes & 0 \\ \divideontimes & 0 & \divideontimes & 0 \\ 0 & \divideontimes & 0 & \divideontimes \\ 0 & \divideontimes & 0 & 0 \end{bmatrix}, \mathbf{A}_\mathcal{G}(a) = \begin{bmatrix} a_{11} & 0 & a_{13} & 0 \\ a_{21} & 0 & a_{23} & 0 \\ 0 & a_{32} & 0 & a_{34} \\ 0 & a_{42} & 0 & 0 \end{bmatrix},$$

where $a = (a_{11}, a_{13}, a_{21}, a_{23}, a_{32}, a_{34}, a_{42})$.

Projections and linear subspaces. We next recall some basic notions regarding projections in linear algebra (see, e.g., [33]). Two vectors $x, y \in \mathbb{R}^n$ are orthogonal if $x^\mathsf{T} y = 0$; the orthogonal complement of $\mathcal{M} \subset \mathbb{R}^n$ is a linear subspace defined as $M^\perp := \{x \in \mathbb{R}^n \ : \ x^\mathsf{T} y = 0, \forall\ y \in \mathcal{M}\}$. Given two subspaces $\mathcal{M}, \mathcal{N} \subseteq \mathbb{R}^n$, the subspace $\mathcal{W} \subseteq \mathbb{R}^n$ is a direct sum of $\mathcal{M}$ and $\mathcal{N}$, denoted $\mathcal{W} = \mathcal{M} \oplus \mathcal{N}$, if $\mathcal{M} \cap \mathcal{N} = \{0\}$, and $\mathcal{M} + \mathcal{N} = \{u + v \ : \ u \in \mathcal{M}, v \in \mathcal{N}\} = \mathcal{W}$. The subspaces $\mathcal{M}, \mathcal{N} \subset \mathbb{R}^n$ are complementary if $\mathcal{M} \oplus \mathcal{N} = \mathbb{R}^n$. A matrix $\Pi \in \mathbb{R}^{n\times n}$ is called a *projection* in $\mathbb{R}^{n\times n}$ if $\Pi^2 = \Pi$. Given complementary subspaces $\mathcal{M}, \mathcal{N} \subset \mathbb{R}^n$, for any $z \in \mathbb{R}^n$ there exists a unique decomposition of the form $z = x + y$, where $x \in \mathcal{M}$, $y \in \mathcal{N}$. The transformation $\Pi_{\mathcal{M},\mathcal{N}}$, defined by $\Pi_{\mathcal{M},\mathcal{N}} z := x$, is called *projection onto $\mathcal{M}$ along $\mathcal{N}$*, and the transformation $\Pi_{\mathcal{N},\mathcal{M}}$ defined by $\Pi_{\mathcal{N},\mathcal{M}} z := y$ is called *projection onto $\mathcal{N}$ along $\mathcal{M}$*. Moreover, the vector x is the projection of z onto $\mathcal{M}$ along $\mathcal{N}$, and the vector y is the projection of z onto $\mathcal{N}$ along $\mathcal{M}$. The projection $\Pi_{\mathcal{M},\mathcal{M}^\perp}$ onto $\mathcal{M}$ along $\mathcal{M}^\perp$ is called *orthogonal projection onto $\mathcal{M}$*. Because the subspace $\mathcal{M}$ uniquely determines $\mathcal{M}^\perp$, in what follows we will denote $\Pi_{\mathcal{M},\mathcal{M}^\perp}$ in compact form as $\Pi_\mathcal{M}$. General projections that are not orthogonal are called *oblique projections*. We recall the following instrumental results.

*Lemma 2.1: (**[33, Thm. 2.11 and Thm. 2.31]**)* If $\Pi \in \mathbb{R}^{n\times n}$ is a projection with $\mathrm{rank}\,(\Pi) = k$, then there exists an invertible matrix $T \in \mathbb{R}^{n\times n}$ such that

$$\Pi = T \begin{bmatrix} I_k & 0 \\ 0 & 0 \end{bmatrix} T^{-1}.$$

Moreover, if Π is an orthogonal projection, then T can be chosen to be an orthogonal matrix, i.e., $TT^\mathsf{T} = I$. □

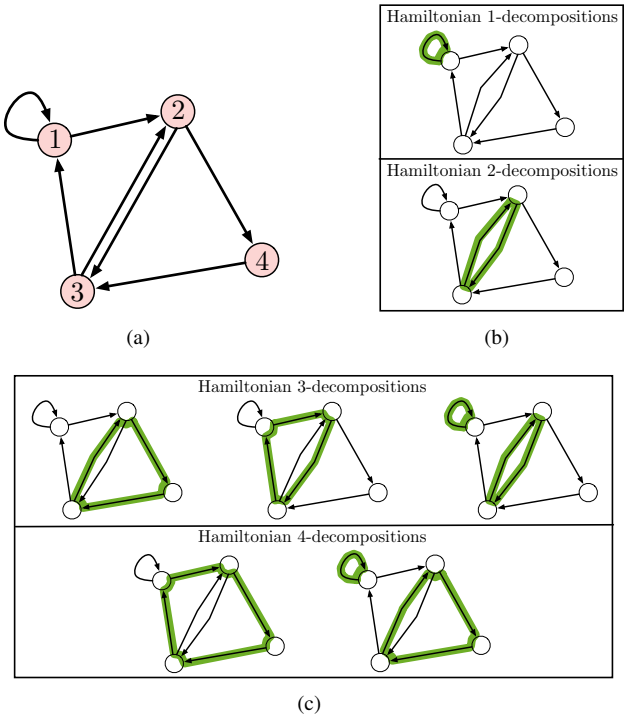

Fig. 2: (a) Illustration of a digraph and (b)-(c) associated Hamiltonian ℓ-decompositions for $\ell \in \{1, \dots, 4\}$. Notice that a graph might admit one, multiple, or no Hamiltonian ℓ-decompositions.

*Lemma 2.2: (**[33, Thm. 2.26]**)* Let $\mathcal{M}, \mathcal{N}$ be complementary subspaces and let the columns of $V \in \mathbb{R}^{n\times k}$ and $W \in \mathbb{R}^{n\times k}$ form a basis for $\mathcal{M}$ and $\mathcal{N}^\perp$, respectively. Then $\Pi_{\mathcal{M},\mathcal{N}} = V(W^\mathsf{T} V)^{-1} W^\mathsf{T}$. □

Recall the following known properties [33, Thm. 1.60]:

$$\mathrm{Im}(M^\mathsf{T}) = \mathrm{Im}(M^\dagger) = \mathrm{Im}(M^\dagger M) = \mathrm{Im}(M^\mathsf{T} M),$$
$$\ker(M) = \mathrm{Im}(M^\mathsf{T})^\perp = \ker(M^\dagger M) = \mathrm{Im}(I - M^\dagger M).$$

From these properties, it follows from Lemma 2.2 that if $M \in \mathbb{R}^{m\times n}$, then $\Pi_{\mathrm{Im}(M)} = MM^\dagger$ and $\Pi_{\ker(M)} = I - M^\dagger M$, where $M^\dagger \in \mathbb{R}^{n\times m}$ is the Moore-Penrose inverse of M.

III. Problem Formulation

Consider n agents whose communication topology is described by a digraph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, $\mathcal{V} = \{1, \dots, n\}$; the state of the i-th agent, $i \in \mathcal{V}$, is $x_i \in \mathbb{R}$. We are interested in a model where each agent exchanges its state with its neighbors and updates it using the following dynamics:

$$\dot{x}_i = a_{ii} x_i + \sum_{j\in\mathcal{N}_i} a_{ij} x_j, \qquad \forall\ i \in \mathcal{V}, \tag{1}$$

where $a_{ij} \in \mathbb{R}$, $(i,j) \in \mathcal{E}$, is a weighting factor. Setting $A = [a_{ij}]$ with $a_{ij} = 0$ if $(i,j) \notin \mathcal{E}$, and $x = (x_1, \dots, x_n)$, the dynamics can be written in vector form as:

$$\dot{x} = Ax. \tag{2}$$

We say that the network dynamics (2) reach a k-dimensional agreement if, asymptotically, each state variable converges to an (agent-dependent) weighted sum of the initial conditions $x(0)$, such that the value at convergence is confined to a subspace of dimension k. This notion is formalized next.

Definition 3.1: ***(k-dimensional agreement)*** Let $k \in \mathbb{N}_{>0}$, $k \leq n-1$, and let $W \in \mathbb{R}^{n\times n}$ satisfy $\operatorname{rank}(W) = k$. The update (2) *enables a k-dimensional agreement on W* if

$$\lim_{t\to\infty} x(t) = Wx(0), \tag{3}$$

for any $x(0) \in \mathbb{R}^n$. □

In this paper, we are interested in cases where the matrix of weights W is fixed but arbitrary; in these cases, we will say that (2) can reach an *agreement on arbitrary weights*.

Definition (3.1) formalizes a notion of "agreement" where, asymptotically, the state of all agents form a vector $x(t)$ that belongs to the k-dimensional space defined by $\operatorname{Im}(W)$. The notion of agreement generalizes the classical notion of consensus; notice, however, that reaching an agreement does not imply that the values of all state variables coincide asymptotically (i.e., it is not guaranteed that $\lim_{t\to\infty} \|x_i(t) - x_j(t)\| = 0$, $i \neq j$). We discuss the relationship between k-dimensional agreement and classical consensus in the following remark.

Remark 3.2: ***(Relationship with consensus problems)*** In the special case where $k = 1$, the matrix W can be written as $W = vw^\mathsf{T}$ for some $v, w \in \mathbb{R}^n$. This corresponds to *scaled consensus*, studied in [29]. When, in addition, $v = \mathbb{1}$ and w is such that $w^\mathsf{T}\mathbb{1} = 1$, we recover the well-studied *consensus* problem [2]. When, additionally, $w = \frac{1}{n}\mathbb{1}$, we recover the *average consensus* problem [2, Sec. 2]. Notice that, $v = \mathbb{1}$ guarantees that all the state variables converge to the same quantity, namely, $\lim_{t\to\infty} \|x_i(t) - x_j(t)\| = 0, \forall\, i, j \in \mathcal{V}$. □

In what follows, we make the following assumption.

Assumption 1: ***(Strong connectivity)*** The communication digraph $\mathcal{G}$ is strongly connected. □

This assumption is motivated by the following remark.

Remark 3.3: ***(Role of strong connectivity)*** When $\mathcal{G}$ is not strongly connected, then for any A that is consistent with $\mathcal{G}$ at least one of the entries of the matrix exponential $\lim_{t\to\infty} e^{At}$ is identically zero (this follows by recalling that $e^{At} = \sum_{i=0}^{\infty} \frac{A^i t^i}{i!}$ and from [34, Cor 4.5]). Thus, when $\mathcal{G}$ is not strongly connected, an agreement cannot be reached on arbitrary weights (notice, however, that agreement might be reached on *some* weights W, with W sparse). □

We conclude this section by illustrating the applicability of the framework in applications in distributed computation.

Example 3.4: ***(Motivating example: distributed computation of k weighted averages)*** Consider a scenario where each agent is interested in computing, using local communication only, an agent-dependent weighted mean (i.e., with agent-dependent weights) of a certain quantity $x^* = (x_1^*, \dots, x_n^*)$, where $x_i^* \in \mathbb{R}$ is known only by agent i. More formally, let $w_i \in \mathbb{R}^n$, with $w_i^\mathsf{T}\mathbb{1} = 1$, denote desired weighting coefficients for agent i, with k of these vectors being linearly independent; we are interested in guaranteeing that the internal state of each agent $x_i(t)$, $i \in \{1, \dots, n\}$, satisfies $\lim_{t\to\infty} x_i(t) = w_i^\mathsf{T} x^*$. This problem emerges commonly in applications in distributed computing where each agent would like to evaluate an agent-dependent function at the common point x^*, whose entries are known only locally. In fact, by denoting by

$$f_i(x) := w_i^\mathsf{T} x, \qquad\qquad i \in \{1, \dots, n\},$$

the objective is to design a distributed protocol of the form (2) that guarantees that $x_i(t) \to f_i(x^*)$ as $t \to \infty$). There are two natural approaches to this problem. The first (and more standard) consists in executing k linear consensus algorithms in parallel [35], where the i-th algorithm is designed to converge asymptotically to $w_i^\mathsf{T} x^*$. A second solution, proposed in this work, consists in executing a single linear algorithm of the form (2). We next detail these two approaches and discuss their communication complexities.

(Approach 1) Consider k duplicates of the agents' states $x^{(1)}, \dots, x^{(k)} \in \mathbb{R}^n$, and let the i-th duplicate implement the protocol $\dot{x}^{(i)} = A^{(i)}x^{(i)}$. A simple choice to achieve $x^{(i)}(t) \to \mathbb{1}w_i^\mathsf{T} x^*$ is to adopt a Laplacian-based consensus algorithm described by $x^{(i)}(0) = x^*$ and $A^{(i)} := P^{(i)} - \operatorname{diag}(P^{(i)}\mathbb{1})$, where $P^{(i)}$ is such that

$$w_i^\mathsf{T}(P^{(i)} - \operatorname{diag}(P^{(i)}\mathbb{1})) = 0.$$

As shown in [35, Thm. 1], this choice guarantees $\lim_{t\to\infty} x^{(i)}(t) = \mathbb{1}w_i^\mathsf{T} x^*$, provided that $\mathcal{G}$ is strongly connected. Unfortunately, it is easy to see that the spatial and communication complexities of this approach do not scale well with the network size. Indeed, each agent must maintain k independent scalar state variables and, at every time step, each agent must transmit these k variables to all its neighbors. It follows that the per-agent spatial complexity is $\mathcal{O}(k)$ (since each agent maintains k copies of a scalar state variable), and the per-agent communication complexity is $\mathcal{O}(k\cdot\deg(\mathcal{G}))$ (and thus is $\mathcal{O}(n \cdot \deg(\mathcal{G}))$ when k is proportional to n). Here, $\deg(\mathcal{G})$ denotes the largest in/out node degree in $\mathcal{G}$.

(Approach 2) To overcome the scalability issue associated with *(Approach 1)*, we propose the use of a linear agreement protocol of the form (2) that achieves agreement as in (3). Here, A is chosen as described shortly in Section V. In this case, the per-agent spatial complexity is $\mathcal{O}(1)$ since agents maintain a single scalar state variable, and the communication complexity is $\mathcal{O}(\deg(\mathcal{G}))$. Fig. 1 provides a comparison of the communication volumes of *(Approach 1)* and *(Approach 2)*. □

IV. Characterization of the Agreement Space and Agreement Algorithms

In this section, we provide a characterization of the agreement space for (2); we also illustrate how to construct a matrix A that defines the agreement protocol. We begin with the following instrumental result.

Lemma 4.1: ***(Spectral properties of agreement protocols)*** The update (2) reaches a k-dimensional agreement if and only if there exists a nonsingular $T \in \mathbb{R}^{n\times n}$ such that A admits the following decomposition:

$$A = T \begin{bmatrix} \mathbb{0} & \mathbb{0} \\ \mathbb{0} & B \end{bmatrix} T^{-1}, \tag{4}$$

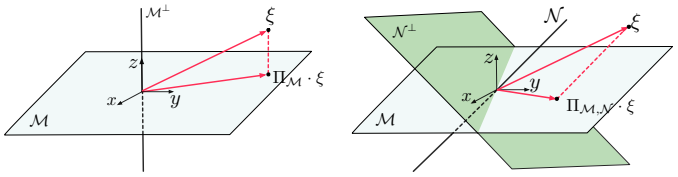

Fig. 3: (a) Illustration of the orthogonal projection of a vector $\xi \in \mathbb{R}^3$ onto $\mathcal{M} \subset \mathbb{R}^3$. (b) Illustration of complementary subspaces $\mathcal{M}, \mathcal{N} \in \mathbb{R}^3$ and oblique projection of ξ. Notice that the projection ray belongs to $\mathrm{span}(\mathcal{N})$.

where $B \in \mathbb{R}^{(n-k)\times(n-k)}$ satisfies $\lambda_{\max}(B) < 0$. □

Proof: *(If)* When (4) holds, we have that:

$$\lim_{t\to\infty} x(t) = \lim_{t\to\infty} e^{At}x(0) = \underbrace{T\begin{bmatrix} I_k & \mathbb{0} \\ \mathbb{0} & \mathbb{0} \end{bmatrix} T^{-1}}_{:=W} x(0) = Wx(0).$$

(Only if) Since $\lim_{t\to\infty} e^{At}$ exists, $\lambda_{\max}(A) \leq 0$; if $\lambda \in \sigma(A)$ and $\Re(\lambda) = 0$, then $\lambda = 0$ and its algebraic and geometric multiplicities coincide (see, e.g., [34, Thm 7.1]). It follows that A must satisfy (4). ■

Lemma 4.1 shows that matrices A that define agreement protocols satisfy two properties: (i) their spectrum is $\sigma(A) = \{0\} \cup \{\lambda_1, \dots, \lambda_{n-k}\}$, where $\Re(\lambda_i) < 0$, and (ii) the eigenvalue $\lambda = 0$ has identical algebraic and geometric multiplicities, which are equal to k.

In the following result, we characterize the class of matrices W on which an agreement can be achieved.

Proposition 4.2: ***(Characterization of agreement space)*** Let $x(t)$ denote the solution of (2) with initial condition $x(0)$. If $\lim_{t\to\infty} x(t) := x_\infty$ exists, then there exist complementary subspaces $\mathcal{M}, \mathcal{N} \subset \mathbb{R}^n$ such that $x_\infty = \Pi_{\mathcal{M},\mathcal{N}} x(0)$. □

Proof: When (2) reaches an agreement, matrix A reads as in (4), from which we obtain

$$W = \lim_{t\to\infty} e^{At} = T\begin{bmatrix} I_k & \mathbb{0} \\ \mathbb{0} & \mathbb{0} \end{bmatrix} T^{-1}.$$

To conclude, notice that W is a projection since $W^2 = W$. ■

Proposition 4.2 provides a characterization of the class of weights on which an agreement can be reached: it asserts that linear protocols of the form (2) can agree on matrix weights that are oblique projections. Accordingly, the asymptotic value at agreement can be interpreted geometrically as the corresponding oblique projection of the initial conditions $x(0)$, cf. Fig. 3 for an illustration.

Remark 4.3: ***(Geometric interpretation of agreement space of consensus algorithms)*** In the case of consensus, the group of agents is known to converge to $\mathbb{1}w^\mathsf{T}x(0)$, where w is a left eigenvector of A ($w^\mathsf{T}A = 0$) that satisfies $w^\mathsf{T}\mathbb{1} = 1$ (see Remark 3.2). Proposition 4.2 provides a geometric interpretation of the consensus value: it shows that $\mathbb{1}w^\mathsf{T}x(0) = \Pi_{\mathcal{M},\mathcal{N}}$, namely, the consensus value is the oblique projection of the initial conditions $x(0)$ onto $\mathcal{M} = \mathrm{Im}(\mathbb{1})$ along $\mathcal{N} = \mathrm{Im}(w)^\perp$ (see Lemma 2.2). In the case of average consensus, the value at convergence is $\frac{1}{n}\mathbb{1}\mathbb{1}^\mathsf{T}x(0)$, which is the orthogonal projection of $x(0)$ onto $\mathcal{M} = \mathrm{Im}(\mathbb{1})$. □

It is now natural to consider the following converse question: given any pair of complementary subspaces $\mathcal{M}, \mathcal{N}$, is it possible to construct a matrix A such that $\lim_{t\to\infty} x(t) = \Pi_{\mathcal{M},\mathcal{N}} x(0)$? The following result provides a positive answer to this question for the case where the digraph $\mathcal{G}$ is complete.

Proposition 4.4: ***(Existence of agreement algorithms over complete digraphs)*** Let $\mathcal{M}, \mathcal{N} \subset \mathbb{R}^n$ be complementary subspaces and $\mathcal{G}$ the complete graph with $|\mathcal{V}| = n$. Then, exists $A \in \mathbb{R}^{n\times n}$ such that the iterates (2) satisfy $\lim_{t\to\infty} x(t) = \Pi_{\mathcal{M},\mathcal{N}} x(0)$. □

Proof: For any pair of complementary subspaces $\mathcal{M}, \mathcal{N}$, [33, Thm. 2.26] guarantees the existence of an oblique projection matrix $\Pi_{\mathcal{M},\mathcal{N}}$. Moreover, by Lemma 2.1, there exists invertible $\bar{T} \in \mathbb{R}^{n\times n}$ such that $\Pi_{\mathcal{M},\mathcal{N}}$ can be decomposed as

$$\Pi_{\mathcal{M},\mathcal{N}} = T_\Pi \begin{bmatrix} I_k & \mathbb{0} \\ \mathbb{0} & \mathbb{0} \end{bmatrix} T_\Pi^{-1},$$

where $k = \dim(\mathcal{M})$. The statement follows by choosing A as in (4) with $T = T_\Pi$ and by noting that, with this choice, $\lim_{t\to\infty} e^{At}x(0) = \Pi_{\mathcal{M},\mathcal{N}} x(0)$. ■

Proposition 4.4 shows that for any pair of complementary subspaces $\mathcal{M}, \mathcal{N}$, it is always possible to construct agreement protocols A such that (2) converges to the projection $\Pi_{\mathcal{M},\mathcal{N}} x(0)$. The proof of this result provides an explicit way to construct A given $\mathcal{M}, \mathcal{N}$. We formalize a technique to construct agreement matrices A by using ideas from Proposition 4.4 in Algorithm 1. We remark that matrices A constructed from Proposition 4.4 or Algorithm 1 are, in general, non-sparse or non-consistent with a pre-specified digraph $\mathcal{G}$, and thus in general they can be interpreted as *centralized* algorithms.

Algorithm 1 Construction of agreement matrix A

Require: $V \in \mathbb{R}^{n\times k}$ whose columns form a basis for $\mathcal{M}$
Require: $U \in \mathbb{R}^{n\times k}$ whose columns form a basis for $\mathcal{N}^\perp$
$\Pi_{\mathcal{M},\mathcal{N}} \leftarrow V(U^\mathsf{T}V)^{-1}U^\mathsf{T}$;
Determine T such that $\Pi_{\mathcal{M},\mathcal{N}} = T\begin{bmatrix} I_k & \mathbb{0} \\ \mathbb{0} & \mathbb{0} \end{bmatrix} T^{-1}$;
Choose $B \in \mathbb{R}^{(n-k)\times(n-k)}$ such that $\lambda_{\max}(B) < 0$;
return $A = T\begin{bmatrix} \mathbb{0}_k & \mathbb{0} \\ \mathbb{0} & B \end{bmatrix} T^{-1}$;

It is thus natural to consider the question of whether a k-dimensional agreement can be achieved by protocols implemented over an arbitrary (non-complete) digraph $\mathcal{G}$ (provided that $\mathcal{G}$ satisfies suitable assumptions). The following example shows that an answer to this question is non-trivial: it shows that if $n = 3$ and $k = 2$, then agreement can be achieved only if $\mathcal{G}$ is the complete graph.

Example 4.5: ***(Not every digraph can reach $k \geq 2$ agreement on arbitrary weights)*** Assume that a network of $n = 3$ agents is interested in agreeing on a space with $k = 2$ by using a sparse communication graph $\mathcal{G}$. By using Lemma 4.1, (2) reaches an agreement if and only if A can be diagonalized as:

$$A = \underbrace{\begin{bmatrix} t_1 & t_2 & t_3 \end{bmatrix}}_{=T} \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \beta \end{bmatrix} \underbrace{\begin{bmatrix} \tau_1 & \tau_2 & \tau_3 \end{bmatrix}^\mathsf{T}}_{=T^{-1}} = \beta t_3 \tau_3^\mathsf{T}, \quad (5)$$

for some β, $\Re(\beta) < 0$, and some $T \in \mathbb{R}^{3\times 3}$ invertible. Since $\mathcal{G}$ is sparse by assumption, at least one of the entries of t_3 or of τ_3 must be equal to zero, which in turn implies that at least one of the rows of A or at least one of the columns of A, respectively, must be identically zero. However, this implies that $\mathcal{G}$ is not strongly connected (see Remark 3.3), and thus at least one of the entries of $W = \lim_{t\to\infty} e^{At}$ must be identically zero. The above discussion shows that if $\mathcal{G}$ is sparse, then it is not possible to achieve an agreement on arbitrary weights W (notice, however, that agreement might be achieved to *some*, graph-dependent, weights W). □

We remark that the above conclusion outlines an emerging behavior with respect to consensus – where strong connectivity of the graph is always sufficient to guarantee existence of a consensus protocol on arbitrary weights – as discussed next.

Remark 4.6: ***(Graphs that guarantee consensus on arbitrary weights)*** We recall that in the case of consensus or, more generally, scaled consensus ($k = 1$), the property of strong connectivity of the underlying digraph is sufficient to guarantee that consensus can be achieved on any arbitrary weighted average (see [35] for consensus or [29] for scaled consensus). More formally, $\forall\ W : \operatorname{rank}(W) = 1, \exists A : \lim_{t\to\infty} x(t) = Wx(0)$. It follows from the discussion in Example 4.5 that agreement on subspaces of dimension $k \geq 2$ mandates the use of digraphs with higher connectivity as compared to the case of consensus. □

Motivated by Example 4.5, in the subsequent section we investigate what properties of (sparse) $\mathcal{G}$ guarantee the existence of agreement protocols that are consistent with $\mathcal{G}$.

V. Agreement algorithms over sparse digraphs

Here, we focus on agreement protocols that are consistent with a given graph $\mathcal{G}$. We adopt the notation $\mathbf{A}_{\mathcal{G}}(a)$ to denote matrices for (2) that are consistent with $\mathcal{G}$ and whose entries are parametrized by the vector $a \in \mathbb{R}^{|\mathcal{E}|}$ (see Section II). Moreover, motivated by Proposition 4.2, we restrict ourselves to cases where W is a projection matrix.

Assumption 2: ***(Matrix of weights is a projection)*** The matrix of weights W is a projection onto a k-dimensional space, namely, $W^2 = W$ and $\operatorname{rank}(W) = k$, with $k \in \mathbb{N}_{>0}$, $k \leq n-1$. □

Building upon Assumption 2, in what follows we make use of the following decomposition for W (see Lemma 2.1):

$$W = T_W \begin{bmatrix} I_k & \mathbb{0} \\ \mathbb{0} & \mathbb{0} \end{bmatrix} T_W^{-1}. \tag{6}$$

Moreover, we often use the following notation:

$$T_W = \begin{bmatrix} t_1 & \cdots & t_n \end{bmatrix}, \quad (T_W^{-1})^{\mathsf{T}} = \begin{bmatrix} \tau_1 & \cdots & \tau_n \end{bmatrix}, \tag{7}$$

where $t_i \in \mathbb{R}^n$, $i \in \{1,\dots,n\}$, denotes the i-th column of T_W, and $\tau_i \in \mathbb{R}^n$ denote the i-th column of T_W and the i-th row of T_W^{-1}, respectively, $i \in \{1,\dots,n\}$.

A. Algebraic conditions for agreement

We next provide an algebraic characterization of agreement protocols over sparse digraphs. Our characterization builds upon a graph-theoretic interpretation of characteristic polynomials presented in [36], which we recall next[1]. In what follows, we denote by $\mathcal{C}_\ell(\mathcal{G})$ the set of all Hamiltonian ℓ-decomposition of the graph $\mathcal{G}$ (see Section II).

Lemma 5.1: ***([36, Thm. 1])*** Let $\mathcal{G}$ be a digraph, let $A \in \mathcal{A}_{\mathcal{G}}$, and $\det(\lambda I - A) = \lambda^n + p_1\lambda^{n-1} + \cdots + p_{n-1}\lambda + p_n$ denote its characteristic polynomial. Each coefficient $p_\ell, \ell \in \{1,\dots,n\}$, can be written as:

$$p_\ell = \sum_{\xi\in\mathcal{C}_\ell(\mathcal{G})} (-1)^{d(\xi)} \prod_{(i,j)\in\xi} a_{ij},$$

where $d(c)$ denotes the number of Hamiltonian cycles in ξ. □

The lemma provides a graph-theoretic description of the characteristic polynomial: it shows that the ℓ-th coefficient of $\det(\lambda I - A)$ is a sum of terms such that each summand is the product of edges in a Hamiltonian ℓ-decomposition of $\mathcal{G}$. We illustrate the claim through an example.

Example 5.2: Consider the digraph in Fig. 2(a). We have:

$$A = \begin{bmatrix} a_{11} & 0 & a_{13} & 0 \\ a_{21} & 0 & a_{23} & 0 \\ 0 & a_{32} & 0 & a_{34} \\ 0 & a_{42} & 0 & 0 \end{bmatrix},$$

and we refer to Fig. 2(b)-(c) for an illustration of all Hamiltonian ℓ-decompositions for this graph. Lemma 5.1 yields:

$$p_1 = -a_{11}, \qquad p_3 = -a_{13}a_{21}a_{32} + a_{11}a_{23}a_{32} - a_{23}a_{42}a_{34},$$
$$p_2 = -a_{23}a_{32} \qquad p_4 = -a_{13}a_{21}a_{42}a_{34} + a_{11}a_{23}a_{34}a_{42}.$$

Notice that each summand in p_ℓ is the product of weights in a Hamiltonian decomposition of the corresponding size. □

We are now ready to state the following necessary and sufficient conditions for agreement.

Theorem 5.3: ***(Algebraic characterization of sparse agreement matrices)*** Let $\mathcal{G}$ be a digraph as in Assumption 1, and let W satisfy Assumption 2. The linear update $\dot{x} = \mathbf{A}_{\mathcal{G}}(a)x$ reaches a k-dimensional agreement on W if and only if the following conditions hold simultaneously:

(i) $\mathbf{A}_{\mathcal{G}}(a)t_i = 0, \quad \tau_i^{\mathsf{T}}\mathbf{A}_{\mathcal{G}}(a) = 0, \qquad \forall i \in \{1,\dots,k\};$

(ii) The polynomial $\lambda^{n-k-1} + p_1\lambda^{n-k-2} + \cdots + p_{n-k-1}$, whose coefficients are defined as

$$p_\ell = \sum_{\xi\in\mathcal{C}_\ell(\mathcal{G})} (-1)^{d(\xi)} \prod_{(i,j)\in\xi} a_{ij}, \quad \ell \in \{1,\dots,n-k\},$$

is stable. □

Proof: *(If)* Let A be any matrix that satisfies *(i)-(ii)*. If A is diagonalizable, then, by letting $T = (t_1, \cdots, t_n)$ be the matrix of its right eigenvectors and $(T^{-1})^{\mathsf{T}} = (\tau_1, \cdots, \tau_n)$ be the matrix of its left eigenvectors, we conclude that A satisfies (4) and thus the linear update reaches an agreement

[1]The claim in [36, Thm. 1] is stated in terms of *cycles* and *cycle families* instead than Hamiltonian cycles and Hamiltonian decompositions. In this work, we used the latter wording, which is more standard and better aligned with the recent literature [30].

on W. If A is not diagonalizable, let T be a similarity transformation such that $T^{-1}AT$ is in Jordan normal form:

$$T^{-1}AT = \begin{bmatrix} J_{\lambda_1} & & & \\ & J_{\lambda_2} & & \\ & & \ddots & \\ & & & J_{\lambda_n} \end{bmatrix}, J_{\lambda_i} = \begin{bmatrix} \lambda_1 & 1 & & \\ & \ddots & \ddots & \\ & & & \lambda_1 \end{bmatrix},$$

From *(i)* we conclude that $\lambda = 0$ is an eigenvalue with algebraic multiplicity k, moreover, since the vectors t_i are linearly independent (see (6)), we conclude that its geometric multiplicity is also equal to k, and thus all Jordan blocks associated with $\lambda = 0$ have dimension 1. Namely, $J_{\lambda_1} = \cdots = J_{\lambda_k} = 0$. By combining this with *(ii)*, we conclude that the characteristic polynomial of A is

$$\det(\lambda I - A) = \lambda^n + p_1\lambda^{n-1} + \cdots + p_{n-k-1}\lambda^{k-1},$$

and, since by assumption such polynomial is stable, we conclude that all remaining eigenvalues $\{\lambda_{k+1}, \ldots, \lambda_n\}$ of A satisfy $\Re(\lambda_i) < 0$. Since all Jordan blocks associated with $\lambda = 0$ have dimension 1 and all the remaining eigenvalues of A are stable, we conclude that A admits the representation (4) and thus the linear update reaches an agreement.

(Only if) We will prove this claim by showing that (4) implies *(i)-(ii)*. To prove that *(i)* holds, we rewrite (4) as

$$T^{-1}AT = \begin{bmatrix} \mathbb{0} & \mathbb{0} \\ \mathbb{0} & B \end{bmatrix},$$

and, by taking the first k columns of the above identity we conclude $At_i = 0$, $i \in \{1, \ldots, k\}$, thus showing that *(i)* holds. To prove that *(ii)* holds, notice that (4) implies that the characteristic polynomial of A is a stable polynomial with k roots at zero. Namely,

$$\begin{aligned} \det(\lambda I - A) &= \lambda^k(\lambda - \lambda_1)(\lambda - \lambda_2)\cdots(\lambda - \lambda_{n-k}) \\ &= \lambda^n + p_1\lambda^{n-1} + \cdots + p_{n-k-1}\lambda^{k-1}, \end{aligned}$$

where $\Re(\lambda_i) < 0, i \in \{1, \ldots, n-k\}$, and $p_j, j \in \{1, \ldots, n-k-1\}$, are nonzero real coefficients. The statement *(ii)* thus follows by applying the graph-theoretic interpretation of the coefficients of the characteristic polynomial in Lemma 5.1. ■

Theorem 5.3 provides an algebraic characterization of linear protocols that reach an agreement on arbitrary weights. Given a graph $\mathcal{G}$ and a matrix of weights W, the result can be used to construct agreement protocols consistent with $\mathcal{G}$ as follows: by interpreting the vector of edge parameters a as an unknown and $p_1, \ldots p_{n-k}$ as pre-specified real numbers, conditions *(i)-(ii)* define a system of $2nk$ linear equations and $n-k$ multilinear polynomial equations with $|\mathcal{E}|$ unknowns described by the vector a and $n-k$ constants $p_1, \ldots p_{n-k}$.

We remark that solvability of *(i)-(ii)* is not guaranteed in general (except for some special cases, such as when the digraph is complete, see Proposition 4.4) and that existence of solutions depends upon: the connectivity of the underlying graph and on the choice of the constants $\{p_1, \ldots p_{n-k-1}\}$. In general cases, solvability can be assessed via standard techniques to solve systems of polynomial equations. We discuss one of these techniques in the following remark.

Remark 5.4: ***(Determining solutions to systems of polynomial equations)*** A powerful and general technique for solving systems of polynomial equations uses the tool of Gröbner bases, as applied using Buchberger's algorithm. The technique relies on transforming a system of polynomial equations into a canonical form, expressed in terms of a Gröbner basis, for which it then easier to determine a solution. We refer to [37], [38] for a complete discussion. Regarding assessing the existence of solutions, we remark that Hilbert's Nullstellensatz theorem provides a powerful tool for this task. In short, the theorem guarantees that a system of polynomial equations has no solution if and only if its Gröbner basis is $\{1\}$. In this sense, the Gröbner basis method provides an easy way to check solvability of *(i)-(ii)*. We also note that the computational complexity of solving a system of polynomial equations via Gröbner bases is exponential [38] in general. □

B. Fast distributed agreement algorithms

The freedom in the choice of the parameters $p_1, \ldots p_{n-k}$ in the statement Theorem 5.3 suggests that a certain graph may admit an infinite number of compatible agreement protocols. For this reason, in this section, we will leverage such freedom to determine agreement protocols with maximal convergence rate. More precisely, we will investigate the following problem: given a digraph $\mathcal{G}$ and a matrix of weights W, determine an agreement protocol that is compatible with $\mathcal{G}$ such that (2) reaches an agreement on W and such that its rate of convergence is maximal. This question can be posed as the following optimization problem:

$$\begin{aligned} \min_{A} \quad & r(A) \\ \text{s.t.} \quad & A \in \mathcal{A}_{\mathcal{G}}, \quad \lim_{t\to\infty} e^{At} = W. \end{aligned} \tag{8}$$

In (8), $r : \mathbb{R}^{n\times n} \to \mathbb{R}$ is a function that measures the rate of convergence of e^{At} over time. The problem of determining the fastest distributed agreement algorithm (8) is closely related to the problem of fastest average consensus studied in [20]; the main difference is that while the average consensus problem is always feasible when $\mathcal{G}$ is strongly connected, there is no simple way to check feasibility of (8) for general graphs. (Notice that feasibility might fail when A is not semi-convergent, and thus $\lim_{t\to\infty} e^{At} = W$ does not exist). When $\mathcal{G}$ is the complete graph Proposition 4.4 guarantees that (8) is feasible.

When the optimization problem (8) is feasible, it is natural to consider two possible choices for the cost function $r(\cdot)$, motivated by the size of $\|e^{At}\|$ as a function of time. The first limiting case is $t \to \infty$. In this case, we consider the following asymptotic measure of convergence motivated by [39, Ch. 14]:

$$r_\infty(A) := \lim_{t\to\infty} t^{-1} \log \|e^{At}\| = \lambda_{\max}(A), \tag{9}$$

where we recall that $\lambda_{\max}(A)$ denotes the spectral abscissa of A (i.e., the largest real part of the eigenvalues of A, see Section II). The second limiting case is $t \to 0$. In this case, we consider the following measure of the initial growth rate

of $\|e^{At}\|$:

$$r_0(A) := \frac{d}{dt}\|e^{At}\|\Big|_{t=0} = \lim_{t\downarrow 0} t^{-1}\log\|e^{At}\| = \lambda_{\max}\left(\frac{A+A^{\mathsf{T}}}{2}\right), \quad (10)$$

where $\lambda_{\max}\left(\frac{A+A^{\mathsf{T}}}{2}\right)$ is the numerical abscissa of A [39]. Accordingly, we have the following two results.

Proposition 5.5: ***(Fast agreement via spectral abscissa minimization)*** Assume that the optimization problem (8) is feasible, and let the performance measure be $r(\cdot) = r_\infty(\cdot)$. Any solution to the following optimization problem:

$$\begin{aligned} \min_{a\in\mathbb{R}^{|\mathcal{E}|}} \quad & \lambda_{\max}\left(\mathbf{A}_{\mathcal{G}}(a)\right) \quad (11)\\ \text{s.t.} \quad & \mathbf{A}_{\mathcal{G}}(a)t_i = 0, \quad \tau_i^{\mathsf{T}}\mathbf{A}_{\mathcal{G}}(a) = 0, \qquad i\in\{1,\dots,k\}, \end{aligned}$$

where t_i, τ_i are as in (7), is also a solution of (8). □

Proof: Since the optimization problem (8) is feasible, condition *(i)* of Theorem 5.3 guarantees that (11) is also feasible and that $\lambda_{\max}\left(\mathbf{A}_{\mathcal{G}}(a)\right) < 0$. Let a^* denote a solution of (11), and let $A = \mathbf{A}_{\mathcal{G}}(a^*)$. By construction, we have $A \in \mathcal{A}_{\mathcal{G}}$, while the two constraints in (11), together with $\lambda_{\max}(A) < 0$ (which is guaranteed by feasibility), guarantee that $\lim_{t\to\infty} e^{At} = W$, which shows that a^* is a feasible point for (8). The claim thus follows by noting that the cost functions of (8) and that of (11) coincide. ■

Proposition 5.5 allows us to recast the optimization problem (8) as a finite-dimensional search over the vector of parameters $a \in \mathbb{R}^{|\mathcal{E}|}$. Unfortunately, even though the constraints of (11) are linear equalities, finding a solution may be computationally burdensome because the objective function (i.e., the spectral abscissa) is not convex (in fact, it is not even Lipschitz [40]). On the other hand, we have the following.

Proposition 5.6: ***(Fast agreement via numerical abscissa minimization)*** Assume that the optimization problem (8) is feasible, and let the performance measure be $r(\cdot) = r_0(\cdot)$. Any solution to the following convex optimization problem:

$$\begin{aligned} \min_{a\in\mathbb{R}^{|\mathcal{E}|}} \quad & \lambda_{\max}\left(\frac{\mathbf{A}_{\mathcal{G}}(a)+\mathbf{A}_{\mathcal{G}}(a)^{\mathsf{T}}}{2}\right) \quad (12)\\ \text{s.t.} \quad & \mathbf{A}_{\mathcal{G}}(a)t_i = 0, \quad \tau_i^{\mathsf{T}}\mathbf{A}_{\mathcal{G}}(a) = 0, \qquad i\in\{1,\dots,k\}, \end{aligned}$$

where t_i, τ_i are as in (7), is also a solution of (8). □

Proof: The proof repeats the same steps as those in the proof of Proposition 5.5, by replacing the cost with (10). ■

In contrast with the spectral abscissa formulation (9), the cost function in (12) is always convex in the parameter a, and thus a global solution of (12) can be computed efficiently.

VI. Graph-theoretic conditions for agreement

Although Theorem 5.3 provides a way to construct agreement protocols that are consistent with a given graph, the system of algebraic equations in its statement may not admit a solution in general (see, e.g., Example 4.5). Motivated by this observation, in this section we provide necessary and sufficient conditions on the graph $\mathcal{G}$ that guarantee solvability of such a set of equations. We begin with a necessary condition.

Proposition 6.1: ***(Graph-theoretic necessary conditions)*** Let $\mathcal{G}$ satisfy Assumption 1. There exists $a \in \mathbb{R}^{|\mathcal{E}|}$ such that the linear update $\dot{x} = \mathbf{A}_{\mathcal{G}}(a)x$ reaches a k-dimensional agreement on arbitrary weights only if

$$|\mathcal{E}| \geq kn. \quad (13)$$

□

Proof: It follows from the algebraic characterization in Theorem 5.3 that $\dot{x} = \mathbf{A}_{\mathcal{G}}(a)x$ reaches an agreement if and only if the following set of algebraic equations admit a solution a:

$$0 = \mathbf{A}_{\mathcal{G}}(a)t_i, \qquad i\in\{1,\dots,k\}, \quad (14a)$$

$$p_\ell = \sum_{\xi\in\mathcal{C}_\ell(\mathcal{G})}(-1)^{d(c)}\prod_{(i,j)\in\xi} a_{ij}, \quad \ell\in\{1,\dots,n-k\}. \quad (14b)$$

The system of equations (14) to be solved consists of nk linearly independent linear equations and $n-k$ nonlinear equations with $|\mathcal{E}|$ unknowns and $n-k$ arbitrarily chosen real numbers $p_1,\dots,p_{n-k}$. Due to the invertibility of matrix T, the equations (14a) are linearly independent, and thus generic solvability of (14) requires the following necessary condition $|a| = |\mathcal{E}| \geq nk$. ■

Condition (13) can be interpreted as a lower bound on the minimal graph connectivity that is required to achieve agreement. The condition shows that the number of edges in $\mathcal{G}$ must grow at least linearly with k or with n. The following remark discusses the above bound when $k = 1$.

Remark 6.2: ***(When $k = 1$, (13) is always guaranteed under standard assumptions)*** By recalling that, when $k = 1$, strong connectivity is a sufficient condition to reach an agreement on arbitrary weights (see [29, Lemma 1]), and by noting that the graph with minimal edge-set cardinality that is strongly connected is the one-directional circular graph (see Fig. 4(a)), it follows that (13) is implicitly guaranteed for standard consensus protocols (i.e. $k = 1$). □

The necessary condition (13) has some important implications for agreement on two widely-studied topologies, namely, line and circulant digraphs; we discuss these cases next.

Remark 6.3: ***(Agreement over circulant digraphs)*** Consider the (one-directional) circulant topology in Fig. 4(a). In this case, $|\mathcal{E}| = 2n$ and thus (13) yields the necessary condition $k \leq 2$. Similarly, consider the (bi-directional) circulant topology in Fig. 4(b). Here, $|\mathcal{E}| = 3n$, and thus (13) yields the necessary condition $k \leq 3$. In words, the one-directional and bi-directional circulant digraphs allow agreement on subspaces of dimension at most 2 and 3, respectively.

It is thus natural to ask the following question: given an arbitrary $k \in \mathbb{N}_{>0}$ and a circulant-type communication digraph where each agent communicates with $\alpha \in \mathbb{N}_{>0}$ neighbors (see Fig. 4(c)), what is the smallest α that is required to reach a k-dimensional agreement? By using (13) with $|\mathcal{E}| = n(\alpha+1)$, an answer to the above question is given by the necessary condition: $\alpha \geq k-1$, which illustrates that the communication degree α must grow at least linearly with k. □

Remark 6.4: ***(Agreement over line digraphs)*** Consider the directed digraph with (bi-directional) line topology illustrated

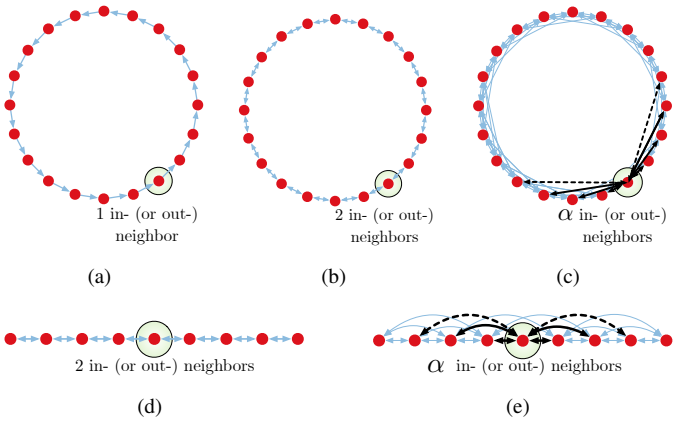

Fig. 4: (a) One-directional circulant topology; (b)–(c) bi-directional circulant topology; (d)-(e) bi-directional line topology. The graph in (a) is the least-connected graph that can reach an agreement on $k = 1$ (see Remark 6.2). (a), (b), and (d) can reach agreement on subspaces of dimension at most $k = 1$ (see Remarks 6.3, 6.4). In (c) and (d), to reach agreement on arbitrary k, α must grow at least linearly with k (see Remarks 6.3, 6.4). In all plots, all nodes have self-cycles, which are omitted here for illustration purposes. Dashed lines illustrate the trend of edge increase as a function of α.

in Fig. 4(d). In this case, $|\mathcal{E}| = n + 2(n-1)$, and thus (17) yields $k \leq \lfloor \frac{3n-2}{n} \rfloor$, which implies $k \leq 3$. Thus, similarly to the circulant digraphs, bi-directional line topologies allow agreement on subspaces of dimension at most 3.

Conversely, it is natural to ask: given any arbitrary $k \in \mathbb{N}_{>0}$ and a line-type communication digraph where each agent communicates with $\alpha \in \mathbb{N}_{>0}$ neighbors (see Fig. 4(e)), what is the smallest α that is required to reach a k-dimensional agreement? By using (13) with $|\mathcal{E}| = n + \alpha n - \frac{\alpha}{2}(\frac{\alpha}{2}+1)$, simple computations yield the necessary condition: $\alpha \geq 2k-1$. Not surprisingly, this condition is more stringent than the circulant topology case (cf. Remark 6.3) since the cardinality of the edge-set of the line topology is smaller than that of the circulant topology (due to the lack of symmetry in the head and tail nodes – see Fig. 4(e)). □

The following result provides a graph-theoretic characterization of graphs that can achieve agreement on arbitrary weights.

Proposition 6.5: ***(Graph-theoretic sufficient conditions)*** Let $\mathcal{G}$ be a graph that satisfies Assumption 1 and let $|\mathcal{E}| \geq nk+n-k$. If there exists a partitioning of the edge parameters $a = (\{a_{ij}\}_{(i,j)\in\mathcal{E}})$ into two disjoint sets $a_v = \{a_1, \dots a_{n-k}\}$ and $a_c = \{a_{n+1}, \dots a_{|\mathcal{E}|}\}$ such that:

- *(i)* For all $\ell \in \{1, \dots, n-k\}$, there exists a Hamiltonian ℓ-decomposition, denoted by $\mathcal{C}_\ell^*$, such that $a_\ell \in \mathcal{C}_\ell^*$;
- *(ii)* Any edge in $\mathcal{C}_\ell^*$ other than a_ℓ belongs to a_c,
- *(iii)* Any Hamiltonian ℓ-decomposition other than $\mathcal{C}_\ell^*$ that contains edges in a_v also contains at least one edge in a_c that does not appear in C_ℓ^*,

then, for any W as in Assumption 2, there exists $a \in \mathbb{R}^{|\mathcal{E}|}$ such that the linear update $\dot{x} = \mathbf{A}_\mathcal{G}(a)x$ reaches a k-dimensional agreement on W. □

Proof: Recall from Thm. 5.3 that $\dot{x} = \mathbf{A}_\mathcal{G}(a)x$ reaches an agreement if and only if there exists a stable polynomial

$$P(\lambda) = \lambda^{n-k-1} + p_1\lambda^{n-k-2} + \dots + p_{n-k-1}. \tag{15}$$

with coefficients $p = (p_1, \dots, p_{n-k})$ such that there exists a solution a^* to the following set of algebraic equations:

$$0 = \mathbf{A}_\mathcal{G}(a)t_i, \qquad i \in \{1, \dots, k\}, \tag{16a}$$

$$p_\ell = \sum_{\xi \in \mathcal{C}_\ell(\mathcal{G})} (-1)^{d(c)} \prod_{(i,j)\in\xi} a_{ij}, \quad \ell \in \{1, \dots, n-k\}. \tag{16b}$$

Hence, we prove this claim by showing that there exists a stable $P(\lambda)$ such that (16) admit a solution. Let $P(\lambda)$ be chosen as follows:

$$P(\lambda) = (\lambda - \alpha_1) \cdots (\lambda - \alpha_n),$$

where its (either real or complex conjugate pairs) roots $\alpha_i \in \mathbb{C}, i \in \{1, \dots, n\}$, satisfy $\Re(\alpha_i) < 0$. Notice that $\Re(\alpha_i) < 0$ imply that all the coefficients $\{p_1, \dots, p_{n-k}\}$ are non-negative. Since $\{\alpha_1, \dots \alpha_n\}$ are arbitrary and for any such choice each element of $p = (p_1, \dots, p_{n-k})$ is non-negative, we will seek solutions of (17) in a neighborhood of $p = 0$.

Since $(t_1, \dots t_k)$ are given (and linearly independent), equation (16a) defines a set of nk linearly independent equations in the variables $a = (a_c, a_v)$, which we denote compactly as $0 = h(a_c, a_v)$, where $h : \mathbb{R}^{|\mathcal{E}|} \to \mathbb{R}^{nk}$. Equation (16b) relates p and (a_c, a_v) by meas of a nonlinear mapping $p = g(a_c, a_v)$, where $g : \mathbb{R}^{|\mathcal{E}|} \to \mathcal{T}$ is a smooth mapping and $\mathcal{T}$ is smooth manifold in $\mathbb{R}^{n-k}$. Since $g(\cdot)$ is a multi-linear polynomial, it is immediate to verify that it admits the following decomposition:

$$g(a_c, a_v) = \frac{\partial g}{\partial a_v} \cdot a_v.$$

By denoting in compact form

$$G(a_c, a_v) := \frac{\partial g}{\partial a_v} \in \mathbb{R}^{n-k \times n-k},$$

$$H(a_c, a_v) := \frac{\partial h}{\partial a_v} \in \mathbb{R}^{nk \times n-k},$$

the system of equations (16) can be rewritten as

$$0 = H(a_c, a_v)a_v, \tag{17a}$$

$$p = G(a_c, a_v)a_v. \tag{17b}$$

As discussed above, we will now seek solutions to (17) in a neighborhood of $p = 0$. By the inverse function theorem [41, Thm. 9.24], solvability of (17) in a neighborhood of $p = 0$ is guaranteed when there exists a particular point (a_c^*, a_v^*) such that $0 = H(a_c^*, a_v^*)a_v^* = G(a_c^*, a_v^*)a_v^*$ and $G(a_c^*, a_v^*)$ is invertible. To show this, we first notice that $a_v^* = 0$ is a solution of (17) for any $a_c \in \mathbb{R}^{|\mathcal{E}|-n+k}$. Thus, we are left to show that there exists a choice a_c^* such that $G(a_c^*, a_v^*)$ is invertible. Thus, we will next provide an inductive method to construct a_c^* such that $G(a_c^*, a_v^*)$ is diagonally dominant.

Let $a_c^{(1)} \in \mathbb{R}^{|\mathcal{E}|-n+k}$ be an arbitrary choice for a_c such that all its entries are nonzero. Notice that condition *(i)* in the statement guarantees that there exists a nonzero product in entry $(1,1)$ of $G(a_c^{(1)}, a_v^*)$, while condition *(ii)* guarantees that such product is independent of a_v^*. Thus, by letting $G^{(1)}(a_c, a_v) := G(a_c, a_v)$, the matrix $G^{(1)}(a_c^{(1)}, a_v^*)$ can be partitioned as:

$$G^{(1)}(a_c^{(1)}, a_v^*) = \begin{bmatrix} G_{11}^{(1)}(a_c^{(1)}) & G_{12}^{(1)}(a_c^{(1)}, a_v^*) \\ G_{21}^{(1)}(a_c^{(1)}, a_v^*) & G_{22}^{(1)}(a_c^{(1)}, a_v^*) \end{bmatrix},$$

where $G_{11}^{(1)} \in \mathbb{R}$, $G_{12}^{(1)} \in \mathbb{R}^{1\times n-k-1}$, $G_{21}^{(1)} \in \mathbb{R}^{n-k-1\times 1}$, $G_{22}^{(1)} \in \mathbb{R}^{n-k-1\times n-k-1}$. By condition *(ii)* and since all entries of $a_c^{(1)}$ are nonzero, we have $G_{11}^{(1)}(a_c^{(1)}) \neq 0$. Moreover, either no element of a_v^* appears in any Hamiltonian 1-decomposition, in which case we have $G_{12}^{(1)}(a_c^{(1)}, a_v^*) = 0$ or, otherwise, by condition *(iii)*, each entry in $G_{12}^{(1)}(a_c^{(1)}, a_v^*)$ is described by a product that contains at least one scalar variable in $a_c^{(1)}$ that does not appear in $G_{11}^{(1)}(a_c^{(1)})$. Denote such scalar variable by $\tilde{a}$ and notice that, by choosing $\tilde{a}$ sufficiently small, the first row of $G^{(1)}$ can be made diagonally dominant. Thus, we update $a_c^{(1)}$ as follows: $a_c^{(2)} = \min\{\tilde{a}, a_c^{(i)}\}$ (where the minimum is taken entrywise).

For the inductive step i, notice that $G^{(i)}(a_c^{(i)}, a_v^*)$ is diagonally dominant if $G_{22}^{(i)}(a_c^{(i)}, a_v^*)$ is diagonally dominant. Thus, by defining $G^{(i+1)}(\cdot,\cdot) = G_{22}^{(i)}(\cdot,\cdot)$, $i \in \{1,\dots,n-1\}$, by letting $a_c^{(i+1)} = \min\{\tilde{a}, a_c^{(i)}\}$ (entrywise minimum), and by iterating the argument, we conclude that $G(a_c^{(n-k)}, a_v^*)$ is diagonally dominant. Invertibility of $G(a_c^*, a_v^*)$ thus follows by letting $a_c^* = a_c^{(n-k)}$, which concludes the proof. ■

Proposition 6.5 provides a set of graph-theoretic properties that are sufficient to guarantee that a certain graph can sustain agreement protocols on arbitrary weights. The result identifies Hamiltonian decompositions as the fundamental component that guarantees the existence of agreement protocols. Indeed, as shown in the proof, the existence of $n-k$ independent Hamiltonian decompositions in $\mathcal{G}$ guarantees that a can be chosen so that $n-k$ modes of $\mathbf{A}(a)$ are stable. Finally, we note that determining the Hamiltonian decompositions of a graph can be done efficiently in polynomial time $\mathcal{O}(n^2)$ by using, e.g., Palmer's Algorithm [42].

The usefulness of Proposition 6.5 depends largely on determining a partitioning of a into two disjoint sets of variables a_v and a_c. An algorithm to determine whether such partitioning exists can be constructed by using ideas similar to [43], where a_v and a_c are derived from a directed spanning tree of $\mathcal{G}$.

Remark 6.6: ***(Minimal graphs for agreement)*** It is worth noting that if $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ admits the set of Hamiltonian ℓ-decompositions $\mathcal{C}_\ell(\mathcal{G})$, then any graph $\mathcal{G}' = (\mathcal{V}', \mathcal{E}')$ such that $\mathcal{V}' = \mathcal{V}$ and $\mathcal{E} \subset \mathcal{E}'$ has a set of Hamiltonian ℓ-decompositions $\mathcal{C}'_\ell(\mathcal{G})$ that satisfies $\mathcal{C}_\ell(\mathcal{G}) \subseteq \mathcal{C}'_\ell(\mathcal{G})$. In other words, any graph obtained by adding edges to $\mathcal{G}$ admits a set of Hamiltonian ℓ-decompositions that includes those of $\mathcal{G}$. It follows that if $\mathcal{G}$ admits agreement protocols on arbitrary weights, then any digraph obtained by adding edges to $\mathcal{G}$ will also admit agreement protocols on arbitrary weights. □

We conclude this section by demonstrating the applicability of Proposition. 6.5 through an example.

Example 6.7: ***(Illustration of Hamiltonian decomposition condition)*** Consider the communication graph illustrated in Fig. 5(a). The corresponding agreement protocol is given by:

$$\mathbf{A}_{\mathcal{G}}(a) = \begin{bmatrix} a_{11} & a_{12} & a_{13} & 0 & a_{15} \\ a_{21} & a_{22} & 0 & 0 & 0 \\ 0 & a_{32} & a_{33} & a_{34} & 0 \\ 0 & 0 & a_{43} & a_{44} & 0 \\ a_{51} & 0 & 0 & a_{54} & a_{55} \end{bmatrix}.$$

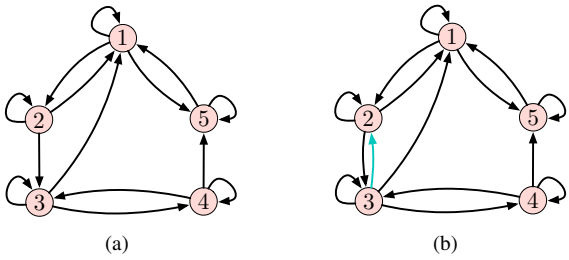

Fig. 5: (a) Example of a graph that can sustain a 2-dimensional agreement. (b) Graph obtained by adding green edges to (a); this graph can sustain a 3-dimensional agreement. See Example 6.7.

By Proposition 6.1, a necessary condition for agreement is

$$k \leq \left\lfloor \frac{|\mathcal{E}|}{n} \right\rfloor = \left\lfloor \frac{14}{5} \right\rfloor = 2,$$

Thus, in what follows we fix $k = 2$. To illustrate the conditions of Proposition 6.5, for simplicity, we let $a_{22} = a_{33} = a_{44} = a_{55} = 0$ (according to Remark 6.6, if the graph without self-cycles has an independent set of Hamiltonian decompositions, then the graph obtained by adding these self-cycles will retain the same set of decompositions). With this choice, the set of all Hamiltonian ℓ-decompositions, $\ell \in \{1,\dots,n-k\}$, is:

$$\begin{aligned} \mathcal{C}_1 &= \{\{a_{11}\}\}, \\ \mathcal{C}_2 &= \{\{a_{12}, a_{21}\}, \{a_{34}, a_{43}\}, \{a_{15}, a_{51}\}\}, \\ \mathcal{C}_3 &= \{\{a_{11}, a_{34}, a_{43}\}, \{a_{13}, a_{21}, a_{32}\}\}. \end{aligned} \tag{18}$$

By selecting a_v and a_c as follows

$$\begin{aligned} a_v &= \{a_{11}, a_{12}, a_{13}\}, \\ a_c &= \{a_{51}, a_{54}, a_{21}, a_{32}, a_{34}, a_{43}, a_{15}\}, \end{aligned}$$

it follows that a set of Hamiltonian ℓ-decompositions that satisfies the conditions in Proposition 6.5 is:

$$\mathcal{C}_1^* = \{a_{11}\}, \quad \mathcal{C}_2^* = \{a_{12}, a_{21}\}, \quad \mathcal{C}_3^* = \{a_{13}, a_{21}, a_{32}\}.$$

Indeed, with this choice, the set of equations (14b) reads as:

$$\begin{bmatrix} p_1 \\ p_2 \\ p_3 \end{bmatrix} = \begin{bmatrix} -1 & 0 & 0 \\ 0 & -a_{21} & 0 \\ a_{34}a_{43} & 0 & -a_{21}a_{32} \end{bmatrix} \begin{bmatrix} a_{11} \\ a_{12} \\ a_{13} \end{bmatrix} - \begin{bmatrix} 0 \\ \gamma \\ 0 \end{bmatrix},$$

where $\gamma = a_{34}a_{43} + a_{15}a_{51}$, which is generically solvable for any $(p_1, p_2, p_3) \in \mathbb{R}^3$. Indeed, any choice of weights such that $a_{21} > 0$ and $|a_{21}a_{32}| > |a_{34}a_{43}|$ guarantees that the above matrix is invertible.

To achieve agreements on subspaces of dimension $k = 3$, consider the graph in Fig. 5(b) obtained by adding edges to the graph in Fig. 5(a). The necessary condition (13) yields

$$k \leq \left\lfloor \frac{|\mathcal{E}|}{n} \right\rfloor = \left\lfloor \frac{15}{5} \right\rfloor = 3,$$

which is satisfied. The set of relevant Hamiltonian decompositions (18) modifies to:

$$\begin{aligned} \mathcal{C}_1 &= \{\{a_{11}\}\}, \\ \mathcal{C}_2 &= \{\{a_{12}, a_{21}\}, \{a_{34}, a_{43}\}, \{a_{15}, a_{51}\}, \{a_{23}, a_{32}\}\}, \end{aligned}$$

By selecting a_v and a_c as follows

$$a_v = \{a_{11}, a_{12}\},$$
$$a_c = \{a_{13}, a_{23}, a_{45}, a_{35}, a_{51}, a_{54}, a_{21}, a_{32}, a_{34}, a_{43}, a_{15}\},$$

a set of Hamiltonian ℓ-decompositions that satisfies Proposition 6.5 is:

$$\mathcal{C}_1^* = \{a_{11}\}, \qquad \mathcal{C}_2^* = \{a_{12}, a_{21}\},$$

thus showing that the sufficient conditions also hold. □

VII. Tracking dynamics for agreement

In analogy with classical consensus processes [44], agreement protocols can be modified to track (the oblique projection of) a time-varying forcing signal $u(t)$. Specifically, given a graph $\mathcal{G}$, consider the network process

$$\dot{x} = \mathbf{A}_{\mathcal{G}}(a)x + \dot{u}, \qquad x(0) = u(0), \tag{19}$$

where a is chosen so that $\mathbf{A}_{\mathcal{G}}(a)$ is an agreement matrix (as in Theorem 5.3) and $u : \mathbb{R}_{\geq 0} \to \mathbb{R}^n$ is a continuously-differentiable function. In this framework, the i-th entry of $\dot{u}$ is known only by agent i, and the objective is to guarantee that $x(t)$ tracks a k-dimensional projection $Wu(t)$ of $u(t)$ asymptotically. The protocol (19) can be interpreted as a generalization of the dynamic average consensus algorithm [44], where the communication matrix is an agreement matrix instead than a Laplacian. The following result characterizes the transient behavior of (19).

Proposition 7.1: ***(Convergence of dynamic agreement protocol)*** Consider the update (19) and let $\mathbf{A}_{\mathcal{G}}(a)$ be an agreement protocol on W as in Theorem 5.3. Then, for all $t \geq 0$:

$$\|x(t) - Wu(t)\| \leq e^{-\hat{\lambda}t}\|x(0) - Wu(0)\| + \frac{1}{\hat{\lambda}} \sup_{0\leq\tau\leq t} \|\dot{u}(\tau)\|, \tag{20}$$

where $\hat{\lambda} = \lambda_{\max}\left(\frac{\mathbf{A}_{\mathcal{G}}(a)+\mathbf{A}_{\mathcal{G}}(a)^{\mathsf{T}}}{2}\right)$. □

Proof: The proof is inspired from [44, Thm. 2] and extends the result to non Laplacian-based protocols and non weight-balanced digraphs. Let W be decomposed as in (6), and consider the following decompositions for T_W and T_W^{-1}:

$$T_W = \begin{bmatrix} T_1 & T_2 \end{bmatrix}, \qquad (T_W^{-1})^{\mathsf{T}} = \begin{bmatrix} U_1 & U_2 \end{bmatrix}, \tag{21}$$

where $T_1, U_1 \in \mathbb{R}^{n\times k}$ and $T_2, U_2 \in \mathbb{R}^{n\times n-k}$. Let $e = x - Wu$ denote the tracking error, and consider the change of variables $\bar{e} = T_W^{-1}e$. In the new variables:

$$\begin{aligned}\dot{\bar{e}} &= T_W^{-1}(\dot{x} - W\dot{u}) \\ &= T_W^{-1}AT_W\bar{e} + T_W^{-1}AWu + T_W^{-1}\dot{u} - T_W^{-1}W\dot{u}, \\ &= T_W^{-1}AT_W\bar{e} + T_W^{-1}\dot{u} - T_W^{-1}W\dot{u},\end{aligned}$$

where the last identity follows by using (6), which implies $AW = 0$. By substituting (21) and by noting that $T^{-1}W = [U_1\ 0]^{\mathsf{T}}$:

$$\begin{aligned}\dot{\bar{e}} &= \begin{bmatrix} U_1^{\mathsf{T}}AT_1 & U_1^{\mathsf{T}}AT_2 \\ U_2^{\mathsf{T}}AT_1 & U_2^{\mathsf{T}}AT_2 \end{bmatrix}\bar{e} + \begin{bmatrix} 0 \\ U_2^{\mathsf{T}} \end{bmatrix}\dot{u} \\ &= \begin{bmatrix} 0 \\ U_2^{\mathsf{T}}AT_2 \end{bmatrix}\bar{e} + \begin{bmatrix} 0 \\ U_2^{\mathsf{T}} \end{bmatrix}\dot{u},\end{aligned} \tag{22}$$

where the last inequality follows by noting that $0 = U_1^{\mathsf{T}}AT_1 = U_1^{\mathsf{T}} = AT_1$ according to condition *(i)* in Theorem 5.3.

Next, decompose $e = (e_1, e_2)$ and $\bar{e} = (\bar{e}_1, \bar{e}_2)$, where $e_1, \bar{e}_1 \in \mathbb{R}^k$ and $e_2, \bar{e}_2 \in \mathbb{R}^{n-k}$, and notice that the following identities hold:

$$\bar{e}_2 = U_2^{\mathsf{T}}e, \qquad e = T_2\bar{e}_2. \tag{23}$$

The first identity follows immediately from (21), while the second follows from (21) and $\bar{e}_1(t) = 0$ at all times. To see that $\bar{e}_1(t) = 0\ \forall t \geq 0$, notice that $\bar{e}_1(0) = U_1^{\mathsf{T}}(x(0) - u(0)) = 0$ thanks to the initialization (19), and $\dot{\bar{e}}_1 = 0$ according to (22). By using (23), we conclude that $\dot{e} = Ae + \dot{u}$, from which (20) follows by noting that

$$e(t) = \exp(At)\cdot e(0) + \int_0^t \exp(A(t-\tau))B\dot{u}(\tau)d\tau,$$

and by using $\|\exp(At)\| \leq \exp\left(-\lambda_{\max}\left(\frac{\mathbf{A}_{\mathcal{G}}(a)+\mathbf{A}_{\mathcal{G}}(a)^{\mathsf{T}}}{2}\right)t\right)$. ■

The error bound (20) shows that the dynamics (19) are input-to-state stable with respect to $\dot{u}$. The bound (20) guarantees that for any forcing signal $u(t)$ with bounded time-derivative the tracking error $\|x(t) - Wu(t)\|$ is bounded at all times. One important implication follows from the statement of the proposition as a subcase: if $\lim_{t\to\infty}\dot{u}(t) = 0$ and thus $\lim_{t\to\infty}u(t) = u^* \in \mathbb{R}^n$, then $\lim_{t\to\infty}x(t) = Wu^*$.

VIII. Applications and numerical validation

In this section, we expand on the distributed computation of averages scenario discussed in Example 3.4 by illustrating our theoretical findings via numerical simulations. We consider two application scenarios.

Applications to sensor measurement de-noising. We consider a problem in distributed computation characterized by a regression model of the form $y = H\theta$, where $H \in \mathbb{R}^{n\times k}, n > k$ and θ is an unknown parameter. We assume that each agent i can sense the i-th entry of vector y, denoted by y_i, and is interested in computing the point $\hat{y}_i$ that is the closest to y_i according to the regression model. To this end, we consider the following regression problem:

$$\theta_{\text{ls}} := \arg\min_{\theta} \|H\theta - y\|. \tag{24}$$

It is well-known that θ_{ls} can be obtained by setting the gradient of $\|H\theta - y\|^2$ equal to zero, which yields $0 = \nabla_\theta\|H\theta - y\|^2 = 2H^{\mathsf{T}}H\theta - 2H^{\mathsf{T}}y$ and thus, when $H^{\mathsf{T}}H$ is invertible, we have $\theta_{\text{ls}} = (H^{\mathsf{T}}H)^{-1}H^{\mathsf{T}}y$. Hence, the desired vector to be computed by the agents (de-noised measurements) is

$$\hat{y} = H\theta_{\text{ls}} = H(H^{\mathsf{T}}H)^{-1}H^{\mathsf{T}}y,$$

which is the orthogonal projection of y onto $\text{Im}(H)$. For figure illustration purposes, in our simulations, we consider the case $n = 50$ (meaning $n = 50$ agents or sensors in the network) and $k = 2$ (meaning the sensor measurements can be interpolated using a line). In our simulations, we computed an agreement protocol by using the optimization problem (12) with agreement weights $W = H(H^{\mathsf{T}}H)^{-1}H^{\mathsf{T}}$ implemented on the circulant graph in Fig. 4(c) with $\alpha = 4$ in/out neighbors. Fig. 6(b) shows the time evolution of the agents states

and Fig. 6(a) shows the sampling points y and asymptotic estimates $\hat{y}$ in comparison with the true regression model. As expected, the distributed algorithm (1) allows the agents to denoise the $n = 50$ sensor measurement, yielding the best linear approximation of the collected data samples. Fig. 6(b) shows the trajectories of the agents states. As expected, at convergence, the states of the agents do not coincide, instead, the agreement state is a 50-dimensional vector constrained on a 2-dimensional space.

Applications to robotic formation control. We next illustrate how agreement protocols can be applied to solve formation control problems [9] in multi-agent robotic networks. Consider a team of $n = 8$ single-integrator robots initially arranged at equal intervals around a unit circle (grey lines in Fig. 7(a)-(c)). By using $x-$ and $y-$coordinates to describe the robots positions, we use $x_0 = (\cos(0), \sin(0), \cos(\frac{\pi}{4}), \sin(\frac{\pi}{4}), \dots, \cos(\frac{7\pi}{4}), \sin(\frac{7\pi}{4})) \in \mathbb{R}^{16}$. Fig. 7 illustrates the trajectories of the robots obtained by using the 2D agreement protocol

$$\dot{x} = (A \otimes I_2)x, \qquad x(0) = x_0,$$

using the circulant communication graph illustrated in Fig. 7(b) with $\alpha = 4$. For comparison, in Fig. 7(a) and (d) we illustrate the trajectories obtained by a consensus algorithm described by weights $W = \frac{1}{n}\mathbb{1}\mathbb{1}^\mathsf{T}$. As expected, the robots meet at the point $(0, 0)$, thus achieving rendezvous [9]. In Fig.7(b) and (e), we illustrate the trajectories resulting from running an agreement protocol (computed by solving (12)) with weights $W = \Pi_{\mathcal{M}}$, where $\Pi_{\mathcal{M}}$ is the orthogonal projection onto $\mathcal{M} = \ker(M_1) \subset \mathbb{R}^8$ with

$$M_1 = \begin{bmatrix} 1 & -1 & -1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & -1 & -1 & 1 \end{bmatrix}.$$

The matrix M_1 encodes attraction and repulsion forces between certain robots at convergence. Indeed, by recalling that the agreement value is $x(\infty) = \Pi_{\mathcal{M}} x(0)$, it follows that at steady state the agents positions satisfy $M_1 x(\infty) = 0$. Hence, the rows of M_1 are interpreted as algebraic constraints on the asymptotic agreement value. From Fig.7(b), which reports the corresponding time-evolution of the $x-$ and $y-$coordinates of the robots, we observe that the robots asymptotically achieve a formation that is characterized by a 2-dimensional subspace. Finally, Fig.s 7(d) and (f) illustrate the trajectories of the robots generated by an agreement protocol (computed by solving (11)) where the weights are described by an oblique projection $W = \Pi_{\mathcal{M},\mathcal{N}}$, where $\mathcal{M} = \ker(M_1) \subset \mathbb{R}^8$ and $\mathcal{N} = \mathrm{Im}(N_1) \subset \mathbb{R}^8$ with

$$N_1^\mathsf{T} = \begin{bmatrix} -1 & 5 & 5 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 5 & 5 & -1 \end{bmatrix}.$$

The use of an oblique projection can be interpreted as a non-homogeneous weighting for the vector that defines the final configuration. Indeed, as shown by the figure, in this case, the robots no longer meet "halfway", instead, robots 2 and 3 [respectively, 6 and 7] travel a longer distance as opposed to robots 1 and 4 [respectively, 5 and 8]).

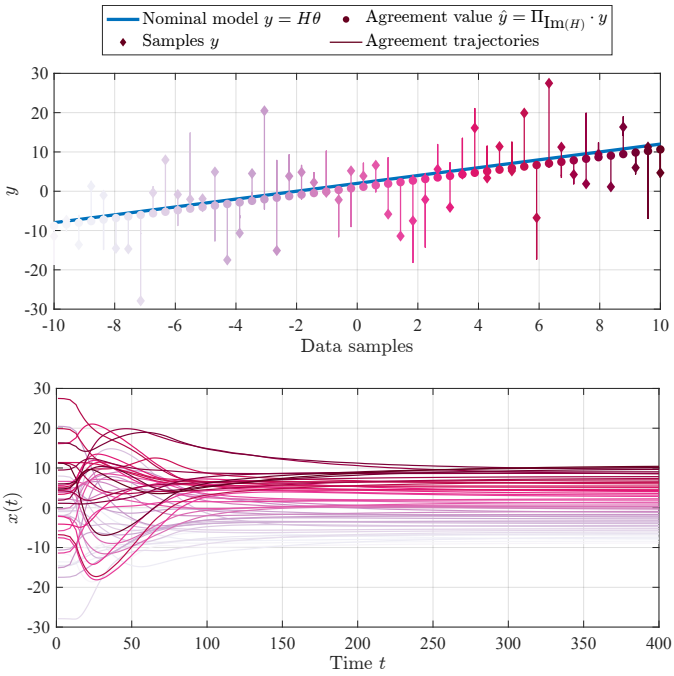

Fig. 6: Application of agreement protocols to solve regression problems. Each agent can measure a sample y_i (represented by diamond markers) and cooperatively computes the projection of $\hat{y}_i$ onto the range of the regression matrix H (represented by circle markers), see (24). In (a), continuous lines illustrate the time evolution of the states of (1). (b) Time evolution of the trajectories of (1).

IX. Conclusions

We have studied the problem of k-dimensional agreement in multi-agent systems, whereby the agents are interested in agreeing on a quantity that belongs to a certain (k-dimensional) subspace without necessarily agreeing on common quantities. We showed that, in general, agreement protocols require communication graphs that are more connected than those needed to achieve consensus, which corresponds to the subcase $k = 1$. To this end, we provided both algebraic and graph-theoretic conditions to identify graphs that can sustain agreement protocols. Our approach provides a graph-theoretic condition that is sufficient to conclude that a certain graph can sustain an agreement protocol but we conjecture this class is much larger in practice. Thus, this work opens the opportunity for several directions for future research, including the possibility of using nonlinear dynamics for agreement, the development of algorithms to compute the agreement weights in a distributed fashion, as well the synthesis of distributed and scalable coordination algorithms to solve optimization problems over networks, where the number of agents and the number of primal variables are of the same order.

References

[1] V. D. Blondel, J. M. Hendrickx, A. Olshevsky, and J. N. Tsitsiklis, "Convergence in multiagent coordination, consensus, and flocking," in *IEEE Conf. on Decision and Control*, 2005, pp. 2996–3000.

[2] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Transactions on Automatic Control*, vol. 49, no. 9, pp. 1520–1533, 2004.

[3] W. Ren, R. W. Beard, and E. M. Atkins, "A survey of consensus problems in multi-agent coordination," in *American Control Conference*, Portland, OR, June 2005, pp. 1859–1864.

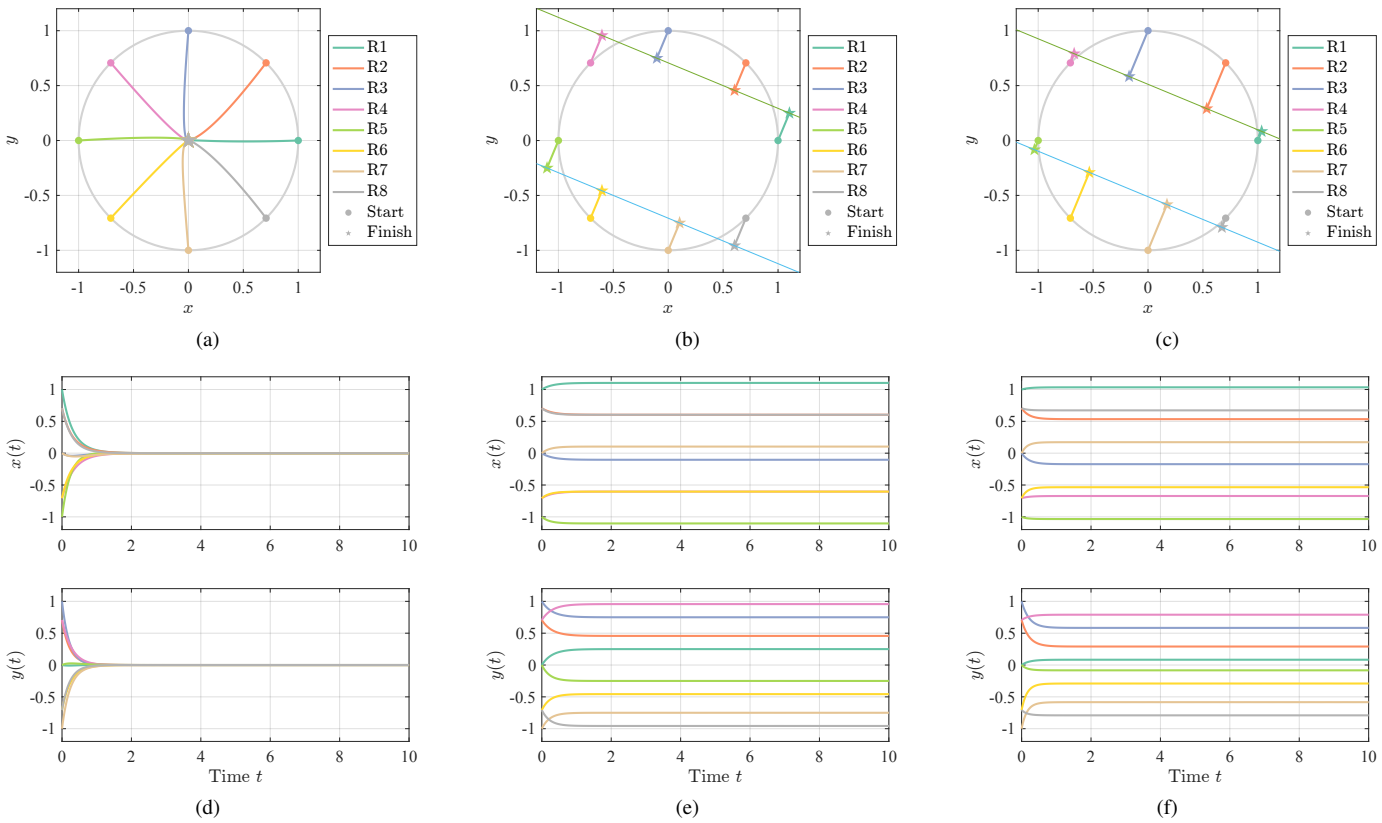

Fig. 7: (a)-(c) Time evolution of the positions of the 8 robots and (d)-(f) trajectories of the $x-$ and $y-$coordinates. (a) and (d) Consensus protocol, which allows the robots to achieve rendezvous. (b) and (e) Agreement protocol on an orthogonal projection onto $\ker(M_1)$. (c) and (f) Agreement on an oblique projection on $\ker(M_1)$ along $\mathrm{Im}(N_1)$.

[4] G. Malewicz, M. H. Austern, A. J.-C. Bik, J. C. Dehnert, I. Horn, N. Leiser, and G. Czajkowski, "Pregel: a system for large-scale graph processing," in *International Conference on Management of data*, 2010, pp. 135–146.

[5] D. Peleg, *Distributed Computing: a Locality-Sensitive Approach*. SIAM, 2000.

[6] 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.

[7] L. Zhao and A. Abur, "Multi area state estimation using synchronized phasor measurements," *IEEE Transactions on Power Systems*, vol. 20, no. 2, pp. 611–617, 2005.

[8] F. Pasqualetti, R. Carli, A. Bicchi, and F. Bullo, "Distributed estimation and detection under local information," in *IFAC Workshop on Distributed Estimation and Control in Networked Systems*, Annecy, France, Sept. 2010, pp. 263–268.

[9] K.-K. Oh, M.-C. Park, and H.-S. Ahn, "A survey of multi-agent formation control," *Automatica*, vol. 53, pp. 424–440, 2015.

[10] J. N. Tsitsiklis, "Problems in decentralized decision making and computation," Ph.D. dissertation, Massachusetts Institute of Technology, Nov. 1984, available at `http://web.mit.edu/jnt/www/Papers/PhD-84-jnt.pdf`.

[11] A. Jadbabaie, J. Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Transactions on Automatic Control*, vol. 48, no. 6, pp. 988–1001, 2003.

[12] M. Cao, A. S. Morse, and B. D. O. Anderson, "Reaching a consensus in a dynamically changing environment - convergence rates, measurement delays and asynchronous events," *SIAM Journal on Control and Optimization*, vol. 47, no. 2, pp. 601–623, 2008.

[13] W. Ren and R. W. Beard, "Consensus seeking in multiagent systems under dynamically changing interaction topologies," *IEEE Transactions on Automatic Control*, vol. 50, no. 5, pp. 655–661, 2005.

[14] J. M. Hendrickx and J. N. Tsitsiklis, "Convergence of type-symmetric and cut-balanced consensus seeking systems," *IEEE Transactions on Automatic Control*, vol. 58, no. 1, pp. 214–218, 2013.

[15] X. Chen, M.-A. Belabbas, and T. Başar, "Consensus with linear objective maps," in *IEEE Conf. on Decision and Control*, 2015, pp. 2847–2852.

[16] I. D. Schizas, A. Ribeiro, and G. B. Giannakis, "Consensus in ad hoc wsns with noisy links—part i: Distributed estimation of deterministic signals," *IEEE Transactions on Signal Processing*, vol. 56, no. 1, pp. 350–364, 2007.

[17] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein, *et al.*, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine learning*, vol. 3, no. 1, pp. 1–122, 2011.

[18] T. Erseghe, D. Zennaro, E. Dall'Anese, and L. Vangelista, "Fast consensus by the alternating direction multipliers method," *IEEE Transactions on Signal Processing*, vol. 59, no. 11, pp. 5523–5537, 2011.

[19] A. Olshevsky and J. N. Tsitsiklis, "Convergence speed in distributed consensus and averaging," *SIAM Journal on Control and Optimization*, vol. 48, no. 1, pp. 33–55, 2009.

[20] L. Xiao and S. Boyd, "Fast linear iterations for distributed averaging," *Systems & Control Letters*, vol. 53, pp. 65–78, 2004.

[21] D. Dolev, N. A. Lynch, S. Pinter, E. W. Stark, and W. E. Weihl, "Reaching approximate agreement in the presence of faults," *Journal of the ACM*, vol. 33, no. 3, pp. 499–516, 1986.

[22] A. Khanafer, B. Touri, and T. Başar, "Robust distributed averaging on networks with adversarial intervention," in *IEEE Conf. on Decision and Control*, 2013, pp. 7131–7136.

[23] A. Nedić, A. Ozdaglar, and P. A. Parrilo, "Constrained consensus and optimization in multi-agent networks," *IEEE Transactions on Automatic Control*, vol. 55, no. 4, pp. 922–938, 2010.

[24] F. Morbidi, "Subspace projectors for state-constrained multi-robot consensus," in *IEEE Int. Conf. on Robotics and Automation*, 2020, pp. 7705–7711.

[25] J. Chen and A. H. Sayed, "Distributed Pareto optimization via diffusion strategies," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 2, pp. 205–220, 2013.

[26] S. Ahmadizadeh, I. Shames, S. Martin, and D. Nešić, "On eigenvalues of Laplacian matrix for a class of directed signed graphs," *Linear Algebra and its Applications*, vol. 523, pp. 281–306, 2017.

[27] W. Li and H. Dai, "Cluster-based distributed consensus," *IEEE Transactions on Wireless Communications*, vol. 8, no. 1, pp. 28–31, 2009.

[28] G. Bianchin, A. Cenedese, M. Luvisotto, and G. Michieletto, "Dis-

tributed fault detection in sensor networks via clustering and consensus," in *IEEE Conf. on Decision and Control*, Osaka, Japan, Dec. 2015, pp. 3828–3833.
[29] S. Roy, "Scaled consensus," *Automatica*, vol. 51, pp. 259–262, 2015.
[30] R. Diestel, *Graph Theory*, 2nd ed., ser. Graduate Texts in Mathematics. Springer, 2000, vol. 173.
[31] C. T. Lin, "Structural controllability," *IEEE Transactions on Automatic Control*, vol. 19, no. 3, pp. 201–208, 1974.
[32] K. Röbenack and K. J. Reinschke, "Digraph based determination of Jordan block size structure of singular matrix pencils," *Linear Algebra and its Applications*, vol. 275-276, pp. 495 – 507, 1998.
[33] A. Galántai, *Projectors and Projection Methods*. Springer Science & Business Media, 2003, vol. 6.
[34] F. Bullo, *Lectures on Network Systems*. Kindle Direct Publishing Santa Barbara, CA, 2019, vol. 1.
[35] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 215–233, Jan 2007.
[36] K. J. Reinschke, "Graph-theoretic approach to symbolic analysis of linear descriptor systems," *Linear Algebra and its Applications*, vol. 197, pp. 217–244, 1994.
[37] D. Cox, J. Little, and D. OShea, *Ideals, Varieties, and Algorithms: an Introduction to Computational Algebraic Geometry and Commutative Algebra*. Springer Science & Business Media, 2013.
[38] M. Kreuzer and L. Robbiano, *Computational Commutative Algebra*. Springer, 2000, vol. 1.
[39] L. N. Trefethen and M. Embree, *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators*. Princeton University Press, 2005.
[40] J. V. Burke and M. L. Overton, "Variational analysis of non-Lipschitz spectral functions," *Mathematical Programming*, vol. 90, no. 2, pp. 317–351, 2001.
[41] W. Rudin, *Principles of Mathematical Analysis*, 3rd ed., ser. International Series in Pure and Applied Mathematics. McGraw-Hill, 1976.
[42] E. M. Palmer, "The hidden algorithm of Ore's theorem on Hamiltonian cycles," *Computers & Mathematics with Applications*, vol. 34, no. 11, pp. 113–119, 1997.
[43] A. Sefik and M. E. Sezer, "Pole assignment problem: a structural investigation," *International Journal of Control*, vol. 54, no. 4, pp. 973–997, 1991.
[44] S. S. Kia, B. Van Scoy, J. Cortés, R. A. Freeman, K. M. Lynch, and S. Martínez, "Tutorial on dynamic average consensus: The problem, its applications, and the algorithms," *IEEE Control Systems Magazine*, vol. 39, no. 3, pp. 40–72, 2019.

Gianluca Bianchin (S'15, M'20) is a Postdoctoral Scholar in the Department of Electrical, Computer, and Energy Engineering at the University of Colorado Boulder. He received the Doctor of Philosophy degree in Mechanical Engineering at the University of California Riverside, in 2020, the Laurea degree in Information Engineering and the Laurea Magistrale degree in Controls Engineering at the University of Padova, Italy, in 2012 and 2014, respectively. He was the recipient of the Dissertation Year Award and of the Dean's Distinguished Fellowship Award from the University of California Riverside.

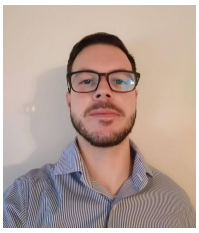

Miguel Vaquero was born in Galicia, Spain. He received his Licenciatura and Master's degree in mathematics from the Universidad de Santiago de Compostela, Spain and the Ph.D. degree in mathematics from Instituto de Ciencias Matemáticas (ICMAT), Spain in 2015. He was then a postdoctoral scholar working on the ERC project "Invariant Submanifolds in Dynamical Systems and PDE" also at ICMAT. From 2017 to 2020, he was a postdoctoral scholar in the Department of Mechanical and Aerospace Engineering of UC San Diego. Since January 2021, he has been an Assistant Professor in the School of Science and Technology at IE University, Madrid, Spain. His interests include optimization, dynamical systems, control theory, machine learning, and geometric mechanics.

Jorge Cortés (M'02, SM'06, F'14) received the Licenciatura degree in mathematics from Universidad de Zaragoza, Zaragoza, Spain, in 1997, and the Ph.D. degree in engineering mathematics from Universidad Carlos III de Madrid, Madrid, Spain, in 2001. He held postdoctoral positions with the University of Twente, Twente, The Netherlands, and the University of Illinois at Urbana-Champaign, Urbana, IL, USA. He was an Assistant Professor with the Department of Applied Mathematics and Statistics, University of California, Santa Cruz, CA, USA, from 2004 to 2007. He is currently a Professor in the Department of Mechanical and Aerospace Engineering, University of California, San Diego, CA, USA. He is a Fellow of IEEE and SIAM. He is the author of Geometric, Control and Numerical Aspects of Nonholonomic Systems (Springer-Verlag, 2002) and co-author (together with F. Bullo and S. Martínez) of Distributed Control of Robotic Networks (Princeton University Press, 2009). At the IEEE Control Systems Society, he has been a Distinguished Lecturer (2010-2014) and an elected member (2018-2020) of its Board of Governors, and is currently its Director of Operations.

Emiliano Dall'Anese is an Assistant Professor in the Department of Electrical, Computer, and Energy Engineering at the University of Colorado Boulder, whee he is also an affiliate Faculty with the Department of Applied Mathematics. He received the Ph.D. in Information Engineering from the Department of Information Engineering, University of Padova, Italy, in 2011. His research interests span the areas of optimization, control, and learning; current applications include power systems and healthcare. He received the National Science Foundation CAREER Award in 2020. He was the recipient of the IEEE PES Prize Paper Award in 2021.