

Consensus Networks over Finite Fields

F. Pasqualetti^a, D. Borra^b, F. Bullo^a

^a*Center for Control, Dynamical Systems, and Computation, University of California, Santa Barbara, USA*

^b*Dipartimento di Scienze Matematiche, Politecnico di Torino, Italy*

Abstract

This work studies consensus strategies for networks of agents with limited memory, computation, and communication capabilities. We assume that agents can process only values from a finite alphabet, and we adopt the framework of finite fields, where the alphabet consists of the integers $\{0, \dots, p-1\}$, for some prime number p , and operations are performed modulo p . Thus, we define a new class of consensus dynamics, which can be exploited in certain applications such as pose estimation in capacity and memory constrained sensor networks. For consensus networks over finite fields, we provide necessary and sufficient conditions on the network topology and weights to ensure convergence. We show that consensus networks over finite fields converge in finite time, a feature that can be hardly achieved over the field of real numbers. For the design of finite-field consensus networks, we propose a general design method, with high computational complexity, and a network composition rule to generate large consensus networks from smaller components. Finally, we discuss the application of finite-field consensus networks to distributed averaging and pose estimation in sensor networks.

1 Introduction

Sensor and actuator networks have recently attracted interest from different research communities, and, in the last years, classic computation, control, and estimation problems have been reformulated to conform the distributed nature of these networked systems [1,2,3]. An important example is the so-called *consensus* problem, where members of a network aim to agree upon a parameter of interest via distributed computation [4]. Consensus algorithms have found broad applicability in many domains, including robotics [5], estimation [6], and parallel computation [7].

In this work we focus on the consensus problem for networks of agents with limited memory, computation, and communication capabilities. We assume that agents are capable of storing, processing, and transmitting exclusively elements from a finite and pre-specified alphabet, and we conveniently model this situation with the formalism of *finite fields*, where the alphabet consists of a set of integers, and operations are performed according to modular arithmetic [8]. We study linear consensus networks over finite fields where, at each time instant, each agent updates its state as a weighted combination over a finite field of its own value and those received from its neighbors. Besides consensus in capacity and memory constrained networks, our finite-field consensus method has broad applicability to problems in cooperative control, networked systems, and network coding, such as averaging, load balancing, and pose estimation from relative measurements. Additionally, the use of a finite alphabet for computation and communication makes our consensus method easily implementable and resilient to communication noise.

Related work Consensus algorithms have been proposed for different network model, agents dynamics, and communication schemes. Starting from the basic setup of time-independent network structure, broadcast and synchronous communication, and unlimited communication bandwidth, consensus algorithms have been proposed to cope, for instance, with time-varying topologies [9], gossip and asynchronous communication [10], and communication errors and link failures [11]. It has been shown that, under mild connectivity assumptions on the interaction graphs, a simple

* This work was supported by NSF grants IIS-0904501 and CNS-1035917, and by ARO grant W911NF-11-1-0092.

Email addresses: `fabiopas@engineering.ucsb.edu` (F. Pasqualetti), `domenica.borra@polito.it` (D. Borra), `bullo@engineering.ucsb.edu` (F. Bullo).

linear iteration ensures consensus, where at each time agents update their state as a weighted average of the available neighboring states [12]. While most of these approaches assume the possibility of processing and transmitting real values, we consider the more realistic case of finite communication bandwidth, possibly due to digital communication channels, and memory constraints. Moreover, we will show that topological conditions ensuring consensus over real values and with real-valued communications are not sufficient for consensus over a finite field.

A relevant body of literature deals with consensus over *quantized* communication channels, where values exchanged by the agents are quantized according to a predefined quantization scheme, and the proposed algorithms are resilient to quantization errors [13,14,15,16,17]. In these works, although the data exchanged among robots is quantized, agents perform computations over the field of real numbers. Thus, the consensus problem with quantized information is related to our problem, yet fundamentally different, since we allow agents to operate on a finite field.

Logical consensus has been studied in [18] for the purpose of intruder or event detection. In logical consensus, a team of agents aims to coordinate their decisions via distributed computation as a function of a set of logical (boolean) events. By leveraging tools from cellular automata and convergence theory of finite-state iteration maps, the focus of [18] is on the design of a synthesis technique for logical consensus systems. The main differences between [18] and this paper are as follows. First, in logical consensus the agents state is a binary variable, while in our work it takes value in an arbitrary finite set. Second, in logical consensus agents are allowed to perform any logical operation, such as {and, or, not}, while our consensus algorithm makes use of only two mathematical operations, namely modular addition and multiplication. Third and finally, in logical consensus agents aim to agree upon a logical expression or compact sets [19], while finite-field consensus algorithms may be used to compute a (non-boolean) function of the initial states, such as the average.

A distributed consensus algorithm with integer communication and computation is proposed and analyzed in [20]. With respect to this work, and similarly to [21], we make use of *modular arithmetic*, instead of standard arithmetic and, therefore, we define a novel and complementary class of consensus networks. The use of modular arithmetic is advantageous in several applications, such as pose estimation from relative measurements (Section 6). Finally, networks based on modular arithmetic are studied in [22], in the context of system controllability and observability, in [23], in the context of (linear) network coding, and in [24], in the more general context of finite dynamical systems.

Contributions The contributions of this paper are fourfold.

First, we propose the use of modular arithmetic to design consensus algorithms for networks of cooperative agents. Consensus networks over finite fields are distributed, require limited, in fact finite, memory, computation, and communication resources, and, as we show, they exhibit finite time convergence. Thus, finite-field consensus algorithms are suitable for capacity and memory constrained networks, and for applications with time constraints.

Second, we thoroughly characterize convergence of consensus networks over finite fields. We provide necessary and sufficient constructive conditions on the network topology and weights to achieve consensus. For instance, we show that a network achieves consensus over a finite field if and only if the network matrix is row-stochastic over the finite field, and its characteristic polynomial is $s^{n-1}(s-1)$. Additionally, we prove that the convergence time of finite-field consensus networks is bounded by the network cardinality, and that graph properties alone are not sufficient to ensure finite-field consensus. Our analysis complements the classic literature on real-valued consensus networks.

Third, we propose systematic methods to design consensus networks over finite fields. In particular, we derive a general design method, with high computational complexity, independent of the agents interaction graph, and a network composition rule based on graph products to generate large consensus networks from smaller components. We show that networks generated by our composition rule exhibit a specific structure, and maintain the convergence properties, including the convergence time, of the underlying components. Moreover, by using our general network design method, we compute finite-field consensus weights for some specific classes of interaction graphs, and we provide a lower bound on the number of networks achieving consensus as a function of the agents interaction graph and the field characteristic.

Fourth and finally, we consider two applications in sensor networks, namely averaging and pose estimation from relative measurements. In the averaging problem agents aim to determine the average (over the field of real numbers) of their initial values. We show that, under a reasonable set of assumptions, the averaging problem can be solved distributively and in finite time by using a finite-field average consensus algorithm. In the pose estimation problem agents aim to estimate their orientation based on (local) relative measurements. For this problem we derive a distributed pose estimation algorithm based on finite-field average consensus, and we characterize its performance.

Paper organization In Section 2 we recall some preliminary notions on fields, linear algebra, and graph theory. Section 3 contains our setup and some preliminary results on consensus over finite fields. In Section 4 we provide necessary and sufficient conditions for consensus over finite fields, and we propose several illustrative examples. In Section 5 we describe methods to design consensus networks over finite fields. Finally, Section 6 contains two application scenarios, and Section 7 concludes the paper.

2 Notation and Preliminary Concepts

In this section we recall some definitions and properties of algebraic fields, linear algebra, and graph theory. We refer the interested reader to [8,25,26] for a comprehensive treatment of these subjects.

A *field* \mathbb{F} is a set of elements together with *addition* and *multiplication* operations, such that the following axioms hold:

- (A1) *Closure* under addition and multiplication, that is, for all $a, b \in \mathbb{F}$, both $a + b \in \mathbb{F}$ and $a \cdot b \in \mathbb{F}$;
- (A2) *Associativity* of addition and multiplication, that is, for all $a, b, c \in \mathbb{F}$, it holds $a + (b + c) = (a + b) + c$ and $a \cdot (b \cdot c) = (a \cdot b) \cdot c$;
- (A3) *Commutativity* of addition and multiplication, that is, for all $a, b \in \mathbb{F}$, it holds $a + b = b + a$ and $a \cdot b = b \cdot a$;
- (A4) *Existence of additive and multiplicative identity* elements, that is, for all $a \in \mathbb{F}$, there exist $b, c \in \mathbb{F}$ such that $a + b = a$ and $a \cdot c = a$;
- (A5) *Existence of additive and multiplicative inverse* elements, that is, for all $a \in \mathbb{F}$, there exist $b, c \in \mathbb{F}$ such that $a + b = 0$ and $a \cdot c = 1$, with $a \neq 0$;
- (A6) *Distributivity* of multiplication over addition, that is, for all $a, b, c \in \mathbb{F}$, it holds $a \cdot (b + c) = (a \cdot b) + (a \cdot c)$.

A field is finite if it contains a finite number of elements. A basic class of finite fields are the fields \mathbb{F}_p with characteristic p a prime number, which consists of the set of integers $\{0, \dots, p - 1\}$, with addition and multiplication defined as in *modular arithmetic*, that is, by performing the operation in the set of integers \mathbb{Z} , dividing by p , and taking the remainder. Unless specified differently, all the operations listed in this paper are performed in the field \mathbb{F}_p .

Let $a : \mathbb{F}_p^m \rightarrow \mathbb{F}_p^n$ be a linear map between the vector spaces of dimensions m and n , respectively, over the field \mathbb{F}_p . As a classical result in linear algebra, the map a can be represented by a matrix A with n rows and m columns, and elements from the field \mathbb{F}_p . The *image* and *kernel* of A are defined as

$$\text{Im}(A) := \{y \in \mathbb{F}_p^n : y = Ax, x \in \mathbb{F}_p^m\}, \quad \text{Ker}(A) := \{x \in \mathbb{F}_p^m : Ax = 0\},$$

where additions and multiplications are performed modulo p . Analogously, the *pre-image* of a set of vectors $V \subseteq \mathbb{F}_p^n$ through A is the set

$$A^{-1}(V) := \{x \in \mathbb{F}_p^m : v = Ax, \text{ for all } v \in V\}.$$

For a matrix $A \in \mathbb{F}_p^{n \times n}$, let $\sigma_p(A)$ denote the set of eigenvalues of A in the field \mathbb{F}_p , that is,

$$\sigma_p(A) = \{\lambda \in \mathbb{F}_p : \det(\lambda I - A) = 0\}.$$

Analogously, let $\mathbb{F}_p[s]$ denote the set of polynomials with coefficients in \mathbb{F}_p , and let $P_A \in \mathbb{F}_p[s]$ denote the characteristic polynomial of A over \mathbb{F}_p .¹ Notice that the cardinality $|\sigma_p(A)|$ may be strictly smaller than the matrix dimension n ; in other words, finite fields are not *algebraically closed*.

We conclude this section with some standard graph definitions. A directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ consists of a set of vertices \mathcal{V} and a set of edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$. An edge $(v, w) \in \mathcal{E}$ is directed from vertex w to vertex v . For a vertex $v \in \mathcal{V}$, the set of in-neighbors of v is defined as $\mathcal{N}_v^{\text{in}} = \{w \in \mathcal{V} : (v, w) \in \mathcal{E}\}$, and the set of out-neighbors as $\mathcal{N}_v^{\text{out}} = \{w \in \mathcal{V} : (w, v) \in \mathcal{E}\}$. The in-degree of $v \in \mathcal{V}$ equals $|\mathcal{N}_v^{\text{in}}|$, whereas the out-degree of $v \in \mathcal{V}$ equals $|\mathcal{N}_v^{\text{out}}|$. A path in G is a subgraph $P = (\{v_1, \dots, v_{k+1}\}, \{e_1, \dots, e_k\})$ such that $v_i \neq v_j$ for all $i \neq j$, and $e_i = (v_i, v_{i+1})$ for

¹ Since the characteristic polynomial $\bar{P}_A(s) \in \mathbb{R}[s]$ contains only integer coefficients for any $A \in \mathbb{F}_p^{n \times n}$, the characteristic polynomial $P_A(s) \in \mathbb{F}_p[s]$ is $\sum_{i=0}^n \text{mod}(\bar{c}_i, p)s^i$, where $\text{mod}(\cdot)$ is the modulus function, and \bar{c}_i is the i -th coefficient of $\bar{P}_A(s)$.

each $i \in \{1, \dots, k\}$. We say that the path starts at v_1 and ends at v_{k+1} , and, at times, we identify a path by its vertex sequence v_1, \dots, v_{k+1} . A cycle or closed path is a path in which the first and last vertex in the sequence are the same, i.e., $v_1 = v_{k+1}$. The length of a path (resp. cycle) equals the number of edges in the path (resp. cycle). A directed graph is strongly (resp. weakly) connected if there exists a directed (resp. undirected) path between any two vertices. Two subgraphs of the same graph are disjoint if they have no common vertices. A *root* (resp. *globally reachable node*) is a vertex v from which (resp. to which) there exists a directed path to (resp. from) every vertex in the graph, including v itself. Finally, a directed graph is aperiodic if there is no integer greater than one that divides the length of every cycle of the graph.

3 Models of Consensus Networks over Finite Fields

Consider a set of $n \in \mathbb{N}_{>0}$ agents and a finite field \mathbb{F}_p , for some prime number p . Let the agents interaction be described by the directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $i \in \mathcal{V}$ denotes the i -th agent, with $\mathcal{V} = \{1, \dots, n\}$, and $(i, j) \in \mathcal{E}$ if there is a directed edge from agent j to agent i (agent i senses agent j , or, equivalently, the behavior of agent j affects agent i). We assume that each agent is able to manipulate and transmit values from the finite field \mathbb{F}_p according to a pre-specified protocol. We focus on distributed protocols in which (i) each agent i is associated with a state $x_i \in \mathbb{F}_p$, and (ii) each agent updates its state as a weighted combination of the states of its in-neighbors $\mathcal{N}_i^{\text{in}}$. Let $a_{ij} \in \mathbb{F}_p$ be the weight associated with the edge (i, j) , and let $A = [a_{ij}]$, $A \in \mathbb{F}_p^{n \times n}$, be the *weighted adjacency matrix* of \mathcal{G} , or simply *network matrix*, where $a_{ij} = 0$ whenever $(i, j) \notin \mathcal{E}$. Let $x : \mathbb{N}_{\geq 0} \rightarrow \mathbb{F}_p^n$ be the vector containing the agents states. Then the evolution of the network state x over time is described by the iteration (or network)

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

where all operations are performed in the field \mathbb{F}_p .

The *transition graph* associated with the iteration (1) over \mathbb{F}_p is defined as $\mathcal{G}_A = (\mathcal{V}_A, \mathcal{E}_A)$, where, $\mathcal{V}_A = \{v : v \in \mathbb{F}_p^n\}$ and, for $v_i, v_j \in \mathcal{V}_A$, the edge $(v_i, v_j) \in \mathcal{E}_A$ if and only if $v_j = Av_i$. It should be observed that the transition graph contains p^n vertices, and that each vertex has unitary out-degree. Moreover, it can be shown that the transition graph is composed of disjoint weakly-connected subgraphs, and that each subgraph contains exactly one cycle, possibly of unitary length [27]. Finally, each disjoint subgraph contains a globally reachable node. This particular structure of the transition graph will be used to derive certain results on finite-field consensus. Examples of transition graphs are given below in Fig. 1 and Fig. 2.

We say that the iteration (1) (or simply the network matrix A) over a finite field achieves

- (i) **asymptotic consensus**, if for all initial states $x(0) \in \mathbb{F}_p^n$ it holds $\lim_{t \rightarrow \infty} x(t) = \alpha \mathbf{1}$, with $\alpha \in \mathbb{F}_p$ and $\mathbf{1} = [1 \dots 1]^T$;
- (ii) **finite-time consensus**, if for all initial states $x(0) \in \mathbb{F}_p^n$ there exists a finite time $T \in \mathbb{N}$ such that $x(T) = x(T + \tau) = \alpha \mathbf{1}$ for all $\tau \in \mathbb{N}$, with $\alpha \in \mathbb{F}_p$ and $\mathbf{1} = [1 \dots 1]^T$.

Consensus networks with real-valued weights and states have been extensively studied [12,28,4]. In this work we show that real-valued consensus networks and finite-field consensus networks exhibit different features, and particular care needs to be taken to ensure the desired properties over finite fields. It is clear from the above definitions that finite-time consensus implies asymptotic consensus. We next show that the converse is also true.

Theorem 3.1 (Asymptotic consensus implies finite-time consensus) *The iteration (1) over the field \mathbb{F}_p achieves asymptotic consensus only if it achieves finite-time consensus.*

PROOF: Let $\mathcal{G}_A = (\mathcal{V}_A, \mathcal{E}_A)$ be the transition graph associated with the iteration (1). Notice that the state trajectory x of (1) coincides with a path on \mathcal{G}_A starting from the vertex $v_0 = x(0)$. Let $\mathcal{C} \subset \mathcal{V}_A$ be the set of consensus vertices, that is, $\mathcal{C} = \{v : v \in \mathcal{V}_A, v = \alpha \mathbf{1}, \alpha \in \mathbb{F}_p\}$. Suppose that the iteration (1) achieves consensus on the value $v_c \in \mathcal{C}$. Since the vertex set \mathcal{V}_A is finite, the distance between v_0 and v_c is also finite. Consequently, a consensus vertex is reached with a path on \mathcal{G} of finite length, that is, a finite number of iterations in (1) are sufficient to achieve consensus. \square

Following Theorem 3.1, iterations over finite fields either achieve consensus in finite time, or they are not convergent.

In view of this result, in what follows we simply use consensus instead of finite-time and asymptotic consensus. Differently from finite-field consensus networks, consensus networks over the field of real numbers usually converge asymptotically. An exception is constituted by the class of de Bruijn graphs, which have been shown to yield finite-time consensus over the field of real numbers [29]. On the other hand, de Bruijn graphs rely on a specific interaction graph, while finite-field consensus networks include a much broader class of interaction graphs. We conclude this section with a simple result. A matrix A over the field \mathbb{F}_p is *nilpotent* if $A^n = 0$ and is *row-stochastic* if $A\mathbf{1} = \mathbf{1}$.

Lemma 3.2 (Finite-field consensus matrices) *Consider the iteration (1) over the field \mathbb{F}_p . If consensus is achieved, then A is either nilpotent or row-stochastic.*

PROOF: Since A achieves consensus, it follows from Theorem 3.1 that $A^t x(0) = A^{t+1} x(0) = \alpha \mathbf{1}$ for some $\alpha \in \mathbb{F}_p$, for all $x(0)$, and for all $t \geq T$, $T \in \mathbb{N}$. Then $A\alpha \mathbf{1} = \alpha \mathbf{1}$, from which we conclude that either $A\mathbf{1} = \mathbf{1}$ (A is row-stochastic) or $\alpha = 0$ for all initial states $x(0)$ (A is nilpotent). \square

As for the case of real-valued consensus, we limit our attention to row-stochastic network matrices. Although consensus is trivially achieved whenever the network matrix is nilpotent, this case is of limited interest because the consensus value is the origin independently of the agents initial states.

4 Analysis of Consensus Networks over Finite Fields

Conditions for consensus in real-valued networks have been deeply investigated in the last years [12,28,4]. For instance, sufficient conditions ensuring real-valued consensus are that the network matrix A is row-stochastic and that the associated directed graph is strongly connected and aperiodic. The following example shows that graph-theoretic properties are not sufficient for an iteration over a finite field to achieve consensus.

Example 1 (Graph properties are not sufficient for finite-field consensus) *Consider a fully connected network with three agents over the field \mathbb{F}_3 . Consider the network matrices*

$$A_1 = \begin{bmatrix} 2 & 1 & 1 \\ 2 & 1 & 1 \\ 2 & 1 & 1 \end{bmatrix}, A_2 = \begin{bmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 2 & 1 \end{bmatrix}, \text{ and } A_3 = \begin{bmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{bmatrix}.$$

Notice that A_1 , A_2 , and A_3 are row-stochastic and their interaction graph is fully connected. It can be verified that over the field \mathbb{F}_3 only the network matrix A_1 achieves consensus, while A_2 and A_3 exhibit oscillatory dynamics for certain initial conditions. An example of oscillatory dynamics generated by A_3 is reported in Table 1.

Table 1
Sample state trajectory for the matrix A_3 in Example 1.

$x(0)$	$x(1)$	$x(2)$	$x(3)$	$x(4)$	$x(5)$	$x(6)$
1	2	0	1	2	0	1
0	1	2	0	1	2	0
0	1	2	0	1	2	0

\square

As shown in Example 1, graph properties of the network matrix are not sufficient to guarantee consensus for iterations over a finite field (in short, *finite-field consensus*). Indeed, although the considered network matrices feature the same connectivity properties, only one of them achieves finite-field consensus. In what follows we provide finite-field consensus conditions based on the algebraic properties of the network matrix, and on the topological properties of its transition graph.

The dynamic behavior of an iteration over a finite field is entirely described by its associated transition graph. The next theorem provides a necessary and sufficient condition for finite-field consensus based on the transition graph.

Theorem 4.1 (Transition graph of a consensus network) *Consider the iteration (1) over the field \mathbb{F}_p with row-stochastic matrix A , and let $\mathcal{G}_A = (\mathcal{V}_A, \mathcal{E}_A)$ be its associated transition graph. The following statements are equivalent:*

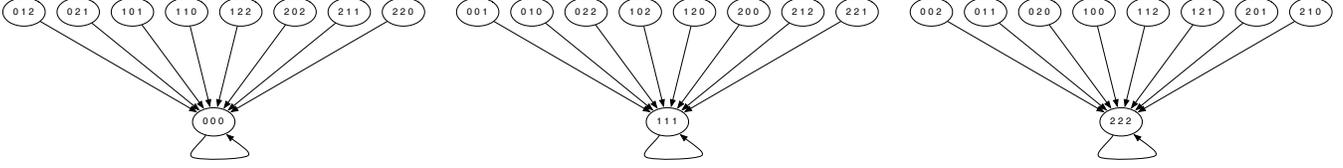


Fig. 1. Transition graph \mathcal{G}_{A_1} associated with the matrix $A_1 \in \mathbb{F}_3^{3 \times 3}$ in Example 1. Since \mathcal{G}_{A_1} contains exactly 3 cycles corresponding to the self-loops around the consensus vertices, the network matrix A_1 achieves consensus (see Theorem 4.1).

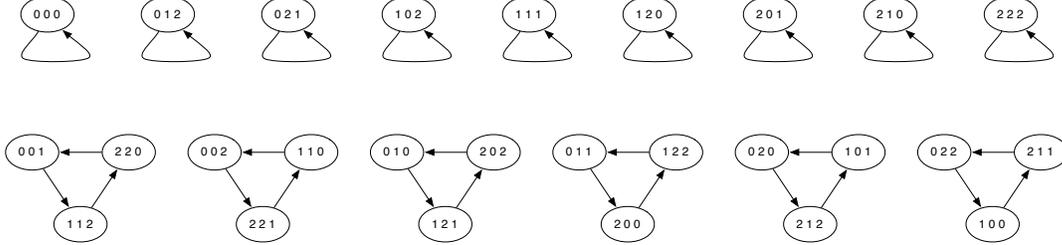


Fig. 2. Transition graph \mathcal{G}_{A_3} associated with the matrix $A_3 \in \mathbb{F}_3^{3 \times 3}$ in Example 1. Since \mathcal{G}_{A_3} contains more than 3 cycles, the network matrix A_3 does not achieve consensus (see Theorem 4.1). The oscillatory state trajectory in Table 1 corresponds to the bottom right cycle in this figure.

- (i) the iteration (1) achieves consensus, and
- (ii) the transition graph \mathcal{G}_A contains exactly p cycles, corresponding to the unitary cycles around the vertices $\alpha \mathbf{1}$, $\alpha \in \{0, \dots, p-1\}$.

Example 2 (Transition graph of a consensus network) The transition graphs associated with the matrices A_1 and A_3 in Example 1 over the field \mathbb{F}_3 are reported in Fig. 1 and Fig. 2, respectively. As previously discussed, and as predicted by Theorem 4.1, the matrix A_1 achieves consensus, while the matrix A_3 does not. \square

PROOF of Theorem 4.1:

(i) \implies (ii) Since the iteration achieves consensus, it follows from Lemma 3.2 that $A\mathbf{1} = \mathbf{1}$. Hence, the transition graph contains p unitary cycles corresponding to the vertices $\alpha \mathbf{1} \in \mathcal{V}_A$, with $\alpha \in \mathbb{F}_p$. Suppose by contradiction that there exists an additional cycle C , and notice that the vertices $\alpha \mathbf{1}$, with $\alpha \in \mathbb{F}_p$, cannot be contained in C since the out-degree of each vertex in the transition graph is exactly one (the transition graph is determined by the linear map A). Thus, there exists a trajectory along C that does not converge to consensus, which contradicts the initial hypothesis.

(ii) \implies (i) Notice that a state trajectory of the iteration (1) is in bijective correspondence with a path on the transition graph \mathcal{G}_A . Suppose that transition graph \mathcal{G}_A contains exactly p unitary cycles located at the vertices $\alpha \mathbf{1} \in \mathcal{V}_A$, with $\alpha \in \mathbb{F}_p$. Then, since each vertex in the transition graph has unitary out-degree, every (sufficiently long) path in \mathcal{G}_A eventually reaches one of the cycles, and, consequently, every state trajectory converges to a consensus state. \square

Theorem 4.1 provides a necessary and sufficient condition for finite-field consensus based on the transition graph. From condition (ii) in Theorem 4.1 and the fact that each vertex in the transition graph has unitary out-degree, we also note that the transition graph of a consensus matrix is composed of p disjoint weakly-connected subgraphs. Moreover, by means of [27, Proposition 3.4], it can be shown that disjoint subgraphs have the same graph topology.

A verification of the convergence condition in Theorem 4.1 may be prohibitive for large networks, because the size of the transition graph grows exponentially with the number of agents in the network (the transition graph contains p^n vertices and p^n edges, since each vertex has unitary out-degree). In what follows we shall derive consensus conditions based on the network matrix instead of its transition graph. Consider the *inverse recursion*

$$\mathcal{S}_\alpha^{t+1} = A^{-1}(\mathcal{S}_\alpha^t), \quad (2)$$

where $\mathcal{S}_\alpha^t \subset \mathbb{F}_p^n$ for all times t and $\mathcal{S}_\alpha^0 = \{\alpha \mathbf{1}\}$. Notice that the inverse recursion defines a sequence of sets, and that the set \mathcal{S}_α^t contains the initial states converging to the consensus value $\alpha \mathbf{1}$ in at most t iterations. We say that the recursion (2) is convergent with limiting set \mathcal{S}_α if there exists $T < n$ satisfying $\mathcal{S}_\alpha = \mathcal{S}_\alpha^T = \mathcal{S}_\alpha^{T+1}$. The following theorem exploits the inverse recursion and the structure of the transition graph to characterize finite-field consensus.

Theorem 4.2 (Recursion subspaces of a consensus network) *Consider the iteration (1) over the field \mathbb{F}_p with row-stochastic matrix A . The following statements are equivalent:*

- (i) the iteration (1) achieves consensus,
- (ii) there exists $\alpha \in \mathbb{F}_p$ such that the recursion (2) is convergent and the limiting set \mathcal{S}_α satisfies $|\mathcal{S}_\alpha| = p^{n-1}$, and
- (iii) for all $\alpha \in \mathbb{F}_p$ the recursion (2) is convergent and each limiting set \mathcal{S}_α satisfies $|\mathcal{S}_\alpha| = p^{n-1}$.

Example 3 (Inverse recursion for a finite-field consensus network) *For the matrix $A_1 \in \mathbb{F}_3^{3 \times 3}$ in Example 1, the set \mathcal{S}_1 generated by the inverse recursion (2) is*

$$\mathcal{S}_1 = \left\{ \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 2 \\ 2 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 2 \end{bmatrix}, \begin{bmatrix} 1 \\ 2 \\ 0 \end{bmatrix}, \begin{bmatrix} 2 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 2 \\ 1 \\ 2 \end{bmatrix}, \begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix} \right\}.$$

Because $|\mathcal{S}_1| = 3^2$, the network matrix A_1 achieves consensus due to Theorem 4.2. Instead, for the network matrix $A_3 \in \mathbb{F}_3^{3 \times 3}$ in Example 1, the inverse recursion yields $\mathcal{S}_1 = \{\mathbf{1}\}$, so that A_3 does not achieve consensus. \square

PROOF of Theorem 4.2: Consider the transition graph $\mathcal{G}_A = (\mathcal{V}, \mathcal{E})$, and define the reverse graph $\bar{\mathcal{G}}_A = (\mathcal{V}, \bar{\mathcal{E}})$, where $(i, j) \in \bar{\mathcal{E}}$ if and only if $(j, i) \in \mathcal{E}$. Notice that the recursion (2) is convergent if and only if $\bar{\mathcal{G}}_A$ contains no cycle of length greater than 1 reachable from $\alpha \mathbf{1}$. Recall from [27, Theorem 1] that \mathcal{G}_A (resp. $\bar{\mathcal{G}}_A$) is obtained as the graph product of a tree by a set of cycles. Hence, the graph \mathcal{G}_A (resp. $\bar{\mathcal{G}}_A$) is composed of disjoint weakly-connected subgraphs, and disjoint subgraphs have the same structure. From this argument we conclude that (ii) and (iii) are equivalent.

(i) \implies (ii) Since A achieves consensus, the graph \mathcal{G}_A contains exactly p unitary cycles corresponding to the consensus vertices (see Theorem 4.1 and Fig. 1). By [27, Theorem 1], the above reasoning, and the fact that A achieves consensus, it follows that $|\mathcal{S}_0| + \dots + |\mathcal{S}_{p-1}| = p^n$, and that $|\mathcal{S}_\alpha| = p^{n-1}$ for all $\alpha \in \mathbb{F}_p$.

(ii) \implies (i) Since $A\mathbf{1} = \mathbf{1}$, the transition graph contains p cycles of unitary length located at the consensus vertices. Let the recursion (2) be convergent for some $\alpha \in \mathbb{F}_p$. From [27, Theorem 1], the graph \mathcal{G}_A contains p identical, disjoint, weakly-connected subgraphs, each one terminating in a consensus vertex. Since $|\mathcal{S}_\alpha| = p^{n-1}$, it follows that consensus is achieved from p^n states (every initial state), which concludes the proof. \square

According to Theorem 4.2, the convergence of the network (1) can be determined by iterating the inverse recursion (2) for some $\alpha \in \mathbb{F}_p$. This computation does not require analyzing the transition graph. Our last and most explicit condition for finite-field consensus is based upon the characteristic polynomial of the network matrix (computed over the finite field).

Theorem 4.3 (Characteristic polynomial of a consensus network) *Consider the iteration (1) over the finite field \mathbb{F}_p with row-stochastic matrix A . The following statements are equivalent:*

- (i) the iteration (1) achieves consensus, and
- (ii) $P_A(s) = s^{n-1}(s - 1)$.

Example 4 (Characteristic polynomial of consensus matrices) *Consider the network matrices in Example 1 over the field \mathbb{F}_3 . It can be verified that*

$$P_{A_1}(s) = s^2(s - 1), \quad P_{A_2}(s) = s(s^2 - 2s + 1), \quad \text{and} \quad P_{A_3}(s) = s^3 - 1.$$

As predicted by our previous analysis and by Theorem 4.3, only the network matrix A_1 achieves consensus. \square

The proof of this theorem is postponed to the Appendix. Theorem 4.3 is equivalently restated as follows: A achieves finite-field consensus if and only if $\sigma_p(A) = \{1, 0, \dots, 0\}$. In other words, the eigenvalues of a finite-field matrix achieving consensus are all contained in the considered finite field and, consequently, every finite-field matrix achieving consensus can be represented in Jordan canonical form via a similarity transformation; see [30] and [31, Theorem 3.5]. We conclude this section by characterizing the convergence value of a finite-field consensus network.

Theorem 4.4 (Finite-field consensus time and value) *Consider the iteration (1) over the finite field \mathbb{F}_p with row-stochastic matrix A and with initial state $x(0) \in \mathbb{F}_p^n$. Assume the iteration (1) achieves consensus. Let $T < n$ denote the dimension of the largest Jordan block associated with the eigenvalue 0. Let $\pi \in \mathbb{F}_p^n$ be the unique eigenvector satisfying $\pi A = \pi$ and $\pi \mathbf{1} = 1$. Then*

$$A^T = \mathbf{1}\pi,$$

so that consensus is achieved at the value $\pi x(0)$ after T iterations. Moreover, the i -th component of π is nonzero only if the i -th vertex of the directed graph associated with A is a root.

PROOF: Since A achieves consensus, we have $\sigma_p(A) = \{1, 0, \dots, 0\}$, and A admits a Jordan canonical form $J_A = V^{-1}AV$ over \mathbb{F}_p . Moreover, the matrix A converges in $T < n$ iterations. The next part of the proof follows the reasoning in [32, Theorem 3]. Let the first column of T be $\mathbf{1}$, and let the matrix J_A have only zero elements, except for the unitary entry in position $(1, 1)$. Since $V^{-1}A = J_A V^{-1}$, the first row of V^{-1} , say π , satisfies $\pi A = \pi$. Then $A^T = V J_A^T V^{-1} = \mathbf{1}\pi$. Since $A\mathbf{1} = \mathbf{1}$, it follows that $\mathbf{1} = A^T \mathbf{1} = \mathbf{1}\pi \mathbf{1}$, and consequently $\pi \mathbf{1} = 1$. To show the last statement, let \mathcal{G} be the directed graph associated with A , and let i be a vertex of \mathcal{G} . Assume that i is not a root of \mathcal{G} , and let the initial state $x(0)$ be all zeros, except for the i -th component. Since i is not a root, there exists a node j which is not reachable from i , and, consequently, the value of the j -th agent is not affected by the i -th agent. Since A achieves consensus for all initial states, the j -th entry of $\mathbf{1}\pi x(0)$ needs to be zero, from which the statement follows. \square

Observe that Theorem 4.4 is not a direct consequence of the theory of non-negative matrices over the field of real numbers [30]. In fact, if regarded as a real-valued matrix, a finite-field consensus matrix is generally unstable.

5 Design of Consensus Networks over Finite Fields

In Section 4 we characterize the convergence of consensus networks over finite fields. With respect to consensus networks over the field of real numbers, finite-field consensus networks require less computational effort and communication bandwidth, and they converge in a finite number of iterations. On the other hand, convergence conditions for finite-field consensus networks depend on the numerical entries of the network matrix (Theorem 4.3), and not only on the connectivity properties of the underlying graph as in the case of consensus networks with real values. For this reason, the design of finite-field consensus networks deserves particular attention. In this section we describe methods to design finite-field consensus networks, and we discuss their limitations. First, we elaborate on Theorem 4.3 to describe a general design method, and we propose a solution for particular interaction graphs. Then we present a scalable composition rule to generate finite-field consensus networks from smaller consensus components.

5.1 Network design via characteristic polynomial

The objective of this section is to design a finite-field consensus matrix A whose sparsity pattern is compatible with a given agents interaction graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, that is, to design the network matrix $A = [a_{ij}]$, $a_{ij} \in \mathbb{F}_p$, where p is a given prime number, and $a_{ij} \neq 0$ only if (i, j) is an edge of \mathcal{G} . Recall from Theorem 4.3 that the network matrix A achieves consensus if and only if its characteristic polynomial is $P_A(s) = s^n(s-1)$ and $A\mathbf{1} = \mathbf{1}$. Since two polynomials are equal if and only if all coefficients are equal, the entries of A can be determined by simultaneously solving the following equations:

$$\begin{cases} A \in \mathbb{F}_p^{n \times n}, \text{ with } a_{ij} = 0 \text{ if } (i, j) \notin \mathcal{E}, \\ \mathbf{1} = A\mathbf{1}, \\ 1 = -c(A, n-1), \\ 0 = c(A, j), \quad j \in \{0, \dots, n-2\}, \end{cases} \quad (3)$$

where $c(A, d)$ is the coefficient of the monomial of degree d in the (parametric) characteristic polynomial P_A .

Example 5 (Design for networks with 2 agents) Let $A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \in \mathbb{F}_p^{2 \times 2}$, and notice that $P_A(s) = s^2 - (a_{11} + a_{22})s + a_{11}a_{22} - a_{12}a_{21}$. It follows from (3) that A achieves consensus if and only if $a_{11} = 1 - a_{12}$, $a_{22} = 1 - a_{21}$, $a_{12} + a_{21} = 1$, and, consequently, if and only if $A = \begin{bmatrix} 1 - \alpha & \alpha \\ 1 - \alpha & \alpha \end{bmatrix}$ for some $\alpha \in \mathbb{F}_p$. \square

Notice that the system of equations (3) contains non-linear, multivariate, polynomial equations, where the unknown variables are the entries of the network matrix A , and that a solution is required over the finite field \mathbb{F}_p . The problem of solving systems of multivariate polynomial equations over finite fields is *NP-hard* [33,34], and it is one of the important research problems in cryptography and information security [35]. Besides enumerating all possible solution candidates,² several solution techniques have been proposed over the last years, see for instance [36,37]. We next provide a condition on the agents interaction graph for the existence of a finite-field network achieving consensus.

Theorem 5.1 (Existence of finite-field consensus matrices) Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be the directed agents interaction graph, with $|\mathcal{V}| = n$ and $|\mathcal{E}| = m$. Assume that \mathcal{G} contains a root, and that $m > \frac{n^2+n}{2}$. Then the system of equations (3) admits $N \geq p^{\frac{2m-n^2-n}{2}}$ solutions over the field \mathbb{F}_p , where N is divisible by the characteristic p . In other words, there exist N network matrices $A = [a_{ij}]$ achieving consensus, with $a_{ij} \in \mathbb{F}_p^{n \times n}$ and $a_{ij} = 0$ if $(i, j) \notin \mathcal{E}$.

PROOF: Let $c(A, d)$ be the coefficient of the monomial of degree d in the parametric characteristic polynomial P_A , with $A = [a_{ij}]$ and $a_{ij} = 0$ if $(i, j) \notin \mathcal{E}$. Define the polynomial $f_i \in \mathbb{F}_p[a_{ij}]$ with $(i, j) \in \mathcal{E}$ as: $f_i = f_i(a_{ij}) = \sum_{j=1}^n a_{ij} - 1$ for $i \in \{1, \dots, n\}$, $f_{n+1} = 1 + c(A, n-1)$, and $f_i = c(A, 2n-i)$ for $i \in \{n+2, \dots, 2n\}$. Notice that a solution to (3), and hence a finite-field consensus matrix compatible with \mathcal{G} , can be computed by simultaneously solving the equations $f_i = 0$ for $i \in \{1, \dots, 2n\}$. Observe that $\deg(f_i) = 1$ for $i \in \{1, \dots, n\}$, and $\deg(f_i) \leq i - n$ for $i \in \{n+1, \dots, 2n\}$. Then $\sum_{i=1}^{2n} \deg(f_i) \leq n + \frac{n(n-1)}{2} \leq \frac{n^2+n}{2}$. Since $m > \frac{n^2+n}{2}$ by assumption, we conclude from [8, Theorem 6.8] that the number of simultaneous solutions N to the equations $f_i = 0$, $i \in \{1, \dots, 2n\}$, in the field \mathbb{F}_p is divisible by the characteristic p . To show that $N \geq p^{\frac{2m-n^2-n}{2}}$, we next construct a network matrix A achieving consensus and compatible with the interaction graph \mathcal{G} , and then employ [8, Theorem 6.11]. Let v be a root of \mathcal{G} , and let $S = (\mathcal{V}, \mathcal{E}_S)$ be a rooted spanning tree of \mathcal{G} with root v , and with $(v, v) \in \mathcal{E}_S$ and $(i, i) \notin \mathcal{E}_S$ for $i \neq v$. Relabel the vertices \mathcal{V} according to their distance from the root v . Define the matrix $A = [a_{ij}]$, where $a_{ij} = 1$ if $(i, j) \in \mathcal{E}_S$, and $a_{ij} = 0$ otherwise. Notice that $A \in \mathbb{F}_p^{n \times n}$ is triangular and row-stochastic, and that its diagonal elements are $\{1, 0, \dots, 0\}$. It follows from Theorem 4.3 that A achieves consensus. \square

In view of the existence Theorem (5.1), consensus over finite fields is possible on a fairly broad class of interaction graphs. A complete characterization of all the interaction topologies yielding consensus over finite fields is beyond the scope of this work, and it is left as the subject of future research. We conclude this section with a remark.

Remark 1 (Network design for fully connected graphs) Let the agents interaction graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be fully connected, that is $(i, j) \in \mathcal{E}$ for all $i, j \in \{1, \dots, n\}$. Let $v \in \mathbb{F}_p^{1 \times n}$ be any vector satisfying $v\mathbf{1} = 1$. Then the network matrix $A = [v^\top \dots v^\top]^\top$ achieves consensus over \mathbb{F}_p ; see for instance the matrix A_1 in Example 1. To see this, let $\mathbf{1}_{\text{orth}} \in \mathbb{F}_p^{n \times (n-1)}$ be any full column rank matrix satisfying $v\mathbf{1}_{\text{orth}} = 0$. Since $A\mathbf{1} = \mathbf{1}$, and $A\mathbf{1}_{\text{orth}} = 0$, we conclude that $\sigma_p(A) = \{1, 0, \dots, 0\}$, with $|\sigma_p(A)| = n$, and the claimed statement follows from Theorem 4.3. \square

5.2 Network design via network composition

In this section we use the concept of graph products to generate finite-field consensus networks from smaller consensus components. For two matrices $A \in \mathbb{F}_p^{n \times n}$ and $B \in \mathbb{F}_p^{m \times m}$, let $A \otimes B \in \mathbb{F}_p^{nm \times nm}$ denote their *Kronecker product*,

² If m is the number of free entries in A , a brute-force solution to (3) requires computing all p^m possible matrices A compatible with the interaction graph \mathcal{G} .

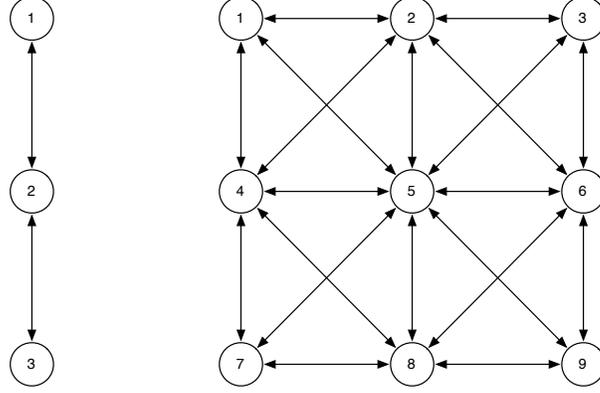


Fig. 3. Agents interaction graphs for the matrices A (left) and $A_k = A \otimes A$ (right) in Example 6. Self-loops have been omitted. Notice that the interaction graphs of A and A_k are self-similar [38].

where [30]

$$A \otimes B = \begin{bmatrix} a_{11}B & a_{12}B & \cdots & a_{1n}B \\ a_{21}B & a_{22}B & \cdots & a_{2n}B \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1}B & a_{n2}B & \cdots & a_{nn}B \end{bmatrix}.$$

Theorem 5.2 (Finite-field consensus via Kronecker product) Consider the network matrices $A \in \mathbb{F}_p^{n \times n}$ and $B \in \mathbb{F}_p^{m \times m}$, and assume that A and B achieve consensus. Then the network matrix $A \otimes B \in \mathbb{F}_p^{nm \times nm}$ achieves consensus.

Example 6 (Finite-field consensus network via Kronecker product) Consider the network matrix

$$A = \begin{bmatrix} 9 & 3 & 0 \\ 1 & 9 & 2 \\ 0 & 7 & 5 \end{bmatrix}$$

over the field \mathbb{F}_{11} . It can be verified that A achieves consensus. By Theorem 5.2 the network matrix

$$A_k = A \otimes A = \begin{bmatrix} 4 & 5 & 0 & 5 & 9 & 0 & 0 & 0 & 0 \\ 9 & 4 & 7 & 3 & 5 & 6 & 0 & 0 & 0 \\ 0 & 8 & 1 & 0 & 10 & 4 & 0 & 0 & 0 \\ 9 & 3 & 0 & 4 & 5 & 0 & 7 & 6 & 0 \\ 1 & 9 & 2 & 9 & 4 & 7 & 2 & 7 & 4 \\ 0 & 7 & 5 & 0 & 8 & 1 & 0 & 3 & 10 \\ 0 & 0 & 0 & 8 & 10 & 0 & 1 & 4 & 0 \\ 0 & 0 & 0 & 7 & 8 & 3 & 5 & 1 & 10 \\ 0 & 0 & 0 & 0 & 5 & 2 & 0 & 2 & 3 \end{bmatrix}$$

achieves consensus over \mathbb{F}_{11} . In fact, it can be verified that $P_{A_k} = s^8(s-1)$ and $A_k \mathbf{1} = \mathbf{1}$. Moreover, the network matrices A and A_k have the same convergence speed. The interaction graph of A and A_k are reported in Fig. 3. \square

PROOF of Theorem 5.2: We first show that $(A \otimes B)$ is row-stochastic. Let $\mathbb{1}_n$ be the vector of all ones of dimension n . Since A and B are row-stochastic over \mathbb{F}_p , we have

$$(A \otimes B)\mathbb{1}_{nm} = \begin{bmatrix} a_{11}B & a_{12}B & \cdots & a_{1n}B \\ a_{21}B & a_{22}B & \cdots & a_{2n}B \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1}B & a_{n2}B & \cdots & a_{nn}B \end{bmatrix} \mathbb{1}_{nm} = \begin{bmatrix} \sum_{j=1}^n a_{1,j}B\mathbb{1}_m \\ \sum_{j=1}^n a_{2,j}B\mathbb{1}_m \\ \vdots \\ \sum_{j=1}^n a_{n,j}B\mathbb{1}_m \end{bmatrix} = \begin{bmatrix} \sum_{j=1}^n a_{1,j}\mathbb{1}_m \\ \sum_{j=1}^n a_{2,j}\mathbb{1}_m \\ \vdots \\ \sum_{j=1}^n a_{n,j}\mathbb{1}_m \end{bmatrix} = \mathbb{1}_{nm}.$$

Following the same argument as in [30], let $J_A = P^{-1}AP$ and $J_B = Q^{-1}BQ$ be the canonical Jordan form of A and B , respectively [31, Theorem 3.5]. Then $J_A \otimes J_B = (P^{-1}AP) \otimes (Q^{-1}BQ) = (P^{-1} \otimes Q^{-1})(A \otimes B)(P \otimes Q) = (P \otimes Q)^{-1}(A \otimes B)(P \otimes Q)$, so that $A \otimes B$ and $J_A \otimes J_B$ are similar. Thus the eigenvalues of $A \otimes B$ are the same as those of $J_A \otimes J_B$, and, because J_A and J_B are upper triangular with eigenvalues λ_i and μ_i on the diagonal, we conclude that $J_A \otimes J_B$ is also upper triangular with diagonal entries, and eigenvalues, $\lambda_i\mu_j$. Finally, since $\sigma_p(A) = \{1, 0, \dots, 0\}$ and $\sigma_p(B) = \{1, 0, \dots, 0\}$, the statement follows from Theorem 4.3. \square

Following Theorem 5.2, finite-field consensus networks can be constructed by composing smaller components. We refer the interested reader to [39,38] for a comprehensive discussion of graphs generated via Kronecker product of adjacency matrices. Regarding the convergence speed of finite-field consensus networks generated via Kronecker products, the following holds. Let $A_k = A_1 \otimes \cdots \otimes A_m$, and let s_i be the number of iterations needed for convergence of the network A_i , $i \in \{1, \dots, m\}$. Due to [40, Theorem 4.3.17], the consensus network A_k converges exactly in $\max\{s_1, \dots, s_m\}$ iterations. In other words, the consensus network A_k is as fast as the slowest of its components.

6 Application to Average Consensus and Distributed Pose Estimation

This section contains two application scenarios for finite-field consensus networks. In Section 6.1 we develop a finite-time averaging algorithm based on finite-field consensus networks. Instead, in Section 6 we use finite-field networks to distributively estimate the orientation of a sensor network given relative measurements.

6.1 Finite-time average consensus

Given a sensor network, let $x_0 \in \mathbb{F}_p^n$ be the vector containing the agents initial states. Let $x_{\mathbb{R}} = \mathbb{1}^T x_0 / n \in \mathbb{R}$ be the average of the agents initial states over the field of real numbers. The average of the agents initial states over the field \mathbb{F}_p can be computed by means of Fermat's little theorem [41] as $x_{\mathbb{F}} = n^{p-2} \mathbb{1}^T x_0 \in \mathbb{F}_p$, where we assume $n \neq kp$, with $k \in \mathbb{N}$, for the inverse of n over \mathbb{F}_p to exist. In what follows, first we show how to compute the average $x_{\mathbb{F}}$ by means of finite-field consensus networks. Then we describe conditions that allow to recover the average $x_{\mathbb{R}}$ from the knowledge of $x_{\mathbb{F}}$ and the total number of agents. We remark that distributed algorithms for average consensus are useful in several applications, including distributed parameter estimation [6], distributed hypothesis testing [42], robotic coordination [2], and for the understanding of opinion dynamics [43].

We say that the iteration (1) over the field \mathbb{F}_p achieves average consensus if it achieves consensus, and the consensus value is $n^{p-2} \mathbb{1}^T x_0$ for every initial state x_0 . A condition for finite-field average consensus is given in the next theorem.

Theorem 6.1 (Finite-field average consensus) Consider the iteration (1) over the field \mathbb{F}_p with row-stochastic matrix A . Assume that the field characteristic satisfies $n \neq kp$ for all $k \in \mathbb{N}$. The following statements are equivalent:

- (i) the iteration (1) achieves average consensus, and
- (ii) $P_A(s) = s^{n-1}(s-1)$, and $\mathbb{1}^T A = \mathbb{1}^T$.

Example 7 (An example of finite-field average consensus) Consider the network matrix

$$A = \begin{bmatrix} 2 & 3 & 1 \\ 2 & 4 & 0 \\ 2 & 4 & 0 \end{bmatrix}$$

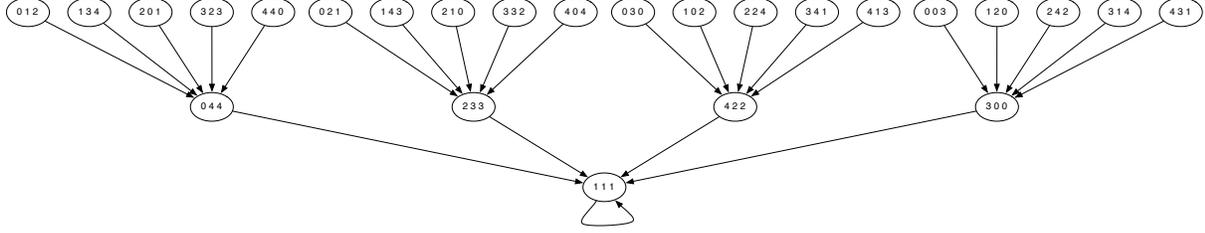


Fig. 4. A subgraph of the transition graph associated with the network matrix A in Example 7. Notice that the sum of the initial states is maintained, and thus average consensus is achieved.

over the field \mathbb{F}_5 . It can be verified that $A\mathbf{1} = \mathbf{1}$, $\mathbf{1}^T A = \mathbf{1}^T$, and $P_A = s^2(s-1)$. By Theorem 6.1 the network matrix A achieves average consensus over \mathbb{F}_5 . In Fig. 4 we show a subgraph of the transition graph associated with A . \square

PROOF of Theorem 6.1:

(i) \implies (ii) Since the iteration achieves consensus, it follows from Theorem 4.4 that $A^n = \mathbf{1}\pi$, where π satisfies $\pi A = \pi$. Because A achieves average consensus, it needs to be $\mathbf{1}\pi = n^{p-2}\mathbf{1}\mathbf{1}^T$. Then $\pi = n^{p-2}\mathbf{1}^T$, and $\mathbf{1}^T A = \mathbf{1}^T$.

(ii) \implies (i) Since $P_A(s) = s^{n-1}(s-1)$ and $A\mathbf{1} = \mathbf{1}$, it follows from Theorem 4.3 that the network achieves consensus. Notice that $\mathbf{1}^T A = \mathbf{1}^T$ implies that $\mathbf{1}^T x(t) = \mathbf{1}^T x(0)$ at all times time t . Let α be the consensus value, and notice that $n\alpha = \mathbf{1}^T x(0)$. To conclude the proof, $\alpha = n^{p-2}\mathbf{1}^T x(0)$, and the network achieves average consensus. \square

Theorem 6.1 provides a necessary and sufficient condition for a network with $n \neq kp$ agents to achieve average consensus over \mathbb{F}_p . The condition $n \neq kp$ is actually necessary for average consensus. In other words, if $n = kp$ for some $k \in \mathbb{N}$, then there exists no network matrix satisfying all conditions in Theorem 6.1 and, therefore, average consensus cannot be achieved. To see this, let $x(0)$ be the network initial state, with $\mathbf{1}^T x(0) \neq 0$, and assume by contradiction that α is the corresponding consensus value. Since $\mathbf{1}^T x(t) = \mathbf{1}^T x(0)$ at all times t , it needs to be $n\alpha = \mathbf{1}^T x(0)$. Then $0 = n\alpha = \mathbf{1}^T x(0) \neq 0$, since kp and 0 are in fact the same element in \mathbb{F}_p .

Suppose now that the average $x_{\mathbb{F}}$ has been computed, and that each agent knows the total number of agents, the field characteristic, and its own initial state. With these assumptions, it is generally not possible to recover the average $x_{\mathbb{R}}$. To see this, consider the case $n = 3$, $p = 5$, and the initial conditions $x_1 = [2 \ 2 \ 2]^T$ and $x_2 = [0 \ 0 \ 1]^T$. Over the field of real numbers we have $x_{1,\mathbb{R}} = \mathbf{1}^T x_1/n = 2$ and $x_{2,\mathbb{R}} = \mathbf{1}^T x_2/n = 1/3$. Over the field \mathbb{F}_p , instead, $x_{1,\mathbb{F}} = n^{p-2}\mathbf{1}^T x_1 = 2$ and $x_{2,\mathbb{F}} = n^{p-2}\mathbf{1}^T x_2 = 2$. Since $x_{1,\mathbb{F}} = x_{2,\mathbb{F}}$ and $x_{1,\mathbb{R}} \neq x_{2,\mathbb{R}}$, it is not possible to recover the average value over the field of real numbers from the average over a finite field and knowledge of network cardinality and parameters. In the next theorem we present a sufficient condition to recover the desired real-valued average.

Theorem 6.2 (Average computation) Let $x_0 \in \mathbb{F}_p^n$, let $x_{\mathbb{R}} = \mathbf{1}^T x_0/n \in \mathbb{R}$, and let $x_{\mathbb{F}} = n^{p-2}\mathbf{1}^T x_0$. If the field characteristic satisfies $n\|x_0\|_{\infty} \leq p$, then $x_{\mathbb{R}} = \text{mod}(n x_{\mathbb{F}}, p)/n$.

PROOF: The statement follows from the relation

$$\text{mod}(n x_{\mathbb{F}}, p) = \text{mod}(n^{p-1}\mathbf{1}^T x_0, p) = \mathbf{1}^T x_0,$$

where the last equality holds because $\text{mod}(n^{p-1}, p) = 1$, and $n\|x_0\|_{\infty} \leq p$. \square

Our finite-time average consensus algorithm follows from Theorems 6.1 and 6.2, and it is reported in Algorithm 1.

We conclude this part by noticing that the condition $n\|x_0\|_{\infty} \leq p$ in Theorem 6.2 is not restrictive. In fact, the field characteristic p is a design parameter, and, in general, it can be chosen to satisfy the above condition as long as the network cardinality n is known and a bound on the agents initial state is known.

Algorithm 1 *Distributed average computation (agent i)*

Input: Field characteristic p , Initial state $x_i(0) \in \mathbb{F}_p$, Neighbors set $\mathcal{N}_i^{\text{in}}$ and $\mathcal{N}_i^{\text{out}}$, Weights $a_{ij} \in \mathbb{F}_p$ for all $j \in \mathcal{N}_i^{\text{in}}$, Number of iterations T (set $T = n$ otherwise), Number of agents n ;

Require: p is a prime number, $A = [a_{ij}]$ achieves average consensus over \mathbb{F}_p , $n \max\{x_1(0), \dots, x_n(0)\} \leq p$;

Output: Average of initial states $x_{\text{ave}} = \sum_{j=1}^n x_j(0)/n \in \mathbb{R}$;

for $t = 0, \dots, T$ **do**

 Transmit $x_i(t)$ to $\mathcal{N}_i^{\text{out}}$;

 Receive $x_j(t)$ from $\mathcal{N}_i^{\text{in}}$;

 Set $x_i(t+1) = \sum_{j \in \mathcal{N}_i^{\text{in}}} a_{ij} x_j(t)$;

end for

return $x_{\text{ave}} = \text{mod}(n x_i(T), p)/n$;

6.2 Pose estimation from relative measurements

In this section we use our previous analysis to calibrate the orientation of a network of cameras. This problem has been previously considered in [44,45,46] as a distributed estimation problem over $SO(2)$. With respect to the existing literature, we let the measurements and the orientations take value in a pre-specified finite field, and we develop an estimation algorithm with performance guarantees based on modular arithmetic.

A camera network is modeled by an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where each vertex is associated with a camera. Let $n = |\mathcal{V}|$ and $m = |\mathcal{E}|$. Let $\theta_i : \mathbb{N}_{\geq 0} \rightarrow \mathcal{O}_p$ be the orientation of the i -th camera as a function of time, where, for some prime number p ,

$$\mathcal{O}_p := \left\{ k \frac{2\pi}{p} : k \in \{0, \dots, p-1\} \right\}. \quad (4)$$

We refer to p as to *discretization accuracy*. For notational convenience, we define the directed graph $\mathcal{G}_d = (\mathcal{V}_d, \mathcal{E}_d)$ associated with the camera network \mathcal{G} , where $\mathcal{V}_d = \mathcal{V}$, and $(i, j) \in \mathcal{E}_d$ if and only if $(i, j) \in \mathcal{E}$ and $i < j$. For each $(i, j) \in \mathcal{E}_d$, let $\eta_{ij} \in \mathcal{O}_p$ be the *relative measurement* between camera i and camera j , that is $\eta_{ij} = \theta_i - \theta_j$. Let θ be the vector of the cameras orientations, and let η be the vector of relative measurements. Assign an arbitrary ordering to the edges \mathcal{E}_d , and define the incidence matrix $B \in \mathbb{F}_p^{m \times n}$ of \mathcal{G}_d by specifying the k -row of B corresponding to the edge (i, j) as

$$b_{kl} = \begin{cases} 1, & \text{if } l = i, \\ -1, & \text{if } l = j, \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

Observe that $B\theta = \eta$. We consider the following problem.

Problem 1 (*Pose estimation from relative measurements*) *Let \mathcal{G} be a camera network, and let \mathcal{G}_d be its associated directed graph. Let p be the discretization accuracy, let B be the incidence matrix of \mathcal{G}_d , and let $\eta \in \mathcal{O}_p^m$ be the vector of relative measurements. Determine a set of cameras orientations $\theta \in \mathcal{O}_p^n$ satisfying $B\theta = \eta$ over \mathbb{F}_p .*

Notice that, if the incidence matrix B and the relative measurements η are available to some camera or central unit, then Problem 1 requires the solution of system of linear equations. Instead, we propose an algorithm that requires each camera to have access to local relative measurements and to communicate with its immediate neighbors. Our distributed pose estimation algorithm is in Algorithm 2.

Because cameras transmit and operate only on values in the finite field \mathbb{F}_p , we argue that Algorithm 2 is suitable for agents with limited capabilities, and it is robust to transmission noise. In the next theorem we analyze the convergence of Algorithm 2.

Theorem 6.3 (*Convergence of Algorithm 2 with perfect measurements*) *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a camera network, and let \mathcal{G}_d be its associated directed graph. Let $\eta \in \mathcal{O}_p^m$ be the vector of relative measurements, and let B be the incidence matrix of \mathcal{G}_d . If $\eta \in \text{Im}(B)$ and A achieves average consensus over \mathbb{F}_p , then*

Algorithm 2 *Distributed pose estimation (camera i)*

Input: Discretization accuracy p , Initial pose $\theta_i(0) \in \mathcal{O}_p$, Neighbors set $\mathcal{N}_i^{\text{in}}$ and $\mathcal{N}_i^{\text{out}}$, Weights $a_{ij} \in \mathbb{F}_p^{n \times n}$ for all $j \in \mathcal{N}_i^{\text{in}}$, Number of iterations T (set $T = n$ otherwise), Relative measurements $\eta_{ij} \in \mathcal{O}_p$ for all $j \in \mathcal{N}_i^{\text{in}}$;

Require: p is a prime number, $A = [a_{ij}]$ achieves average consensus over \mathbb{F}_p ;

Output: Orientation θ_i compatible with measurements $\eta = [\eta_{ij}]$;

for $t = 0, \dots, T$ **do**

 Transmit $x_i(t) = \frac{p\theta_i(t)}{2\pi}$ to $\mathcal{N}_i^{\text{out}}$;

 Receive $x_j(t) = \frac{p\theta_j(t)}{2\pi}$ from $\mathcal{N}_i^{\text{in}}$;

 Update orientation $\theta_i(t)$ as:

$$x_i(t+1) = \sum_{j \in \mathcal{N}_i^{\text{in}}} a_{ij} \left(x_j(t) + \frac{p\eta_{ij}}{2\pi} \right), \quad (6a)$$

$$\theta_i(t+1) = x_i(t+1) \frac{2\pi}{p}. \quad (6b)$$

end for

return Orientation θ_i ;

- (i) Algorithm 2 converges in finite time, that is, $\tilde{\theta} := \theta(T) = \theta(T + \tau)$ for some $\tilde{\theta} \in \mathcal{O}_p^n$, $T < n$, and for all $\tau \in \mathbb{N}$, and
(ii) the final network orientation is compatible with the relative measurements, that is, $B\tilde{\theta} = \eta$.

PROOF: Consider the update law (6a), and notice that it can be written as $x(t+1) = Ax(t) + LBv$, where $\frac{p\eta}{2\pi} = Bv$ for some vector $v \in \mathbb{F}_p^n$ ($y \in \text{Im}(B)$ by assumption), $L \in \mathbb{F}_p^{n \times m}$, and the k -th column of L corresponding to the edge $(i, j) \in \mathcal{E}_d$ is specified as

$$l_{\ell k} = \begin{cases} a_{ij}, & \text{if } \ell = i, \\ -a_{ij}, & \text{if } \ell = j, \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

Observe that

$$(LB)_{ij} = \begin{cases} \sum_{k \in \mathcal{N}_i^{\text{in}}} a_{ik}, & \text{if } i = j, \\ -a_{ij}, & \text{if } j \in \mathcal{N}_i^{\text{in}}, \\ 0, & \text{otherwise.} \end{cases}$$

so that $LB\mathbf{1} = 0$ (asymmetric Laplacian matrix of \mathcal{G} [2]). Notice that $x(t) = A^t x(0) + \sum_{\tau=0}^{t-1} A^\tau LBv$. Since A achieves average consensus, we have $A^T = n^{p-2} \mathbf{1}\mathbf{1}^\top$ for some $T < n$. Thus, $A^t LB = n^{p-2} \mathbf{1}\mathbf{1}^\top LB$ for all $t \geq T$. We now show that $\mathbf{1}^\top LB = 0$, from which statement (i) follows. Since A achieves average consensus, it follows from Theorem 6.1 that $A\mathbf{1} = \mathbf{1}$ and $\mathbf{1}^\top A = \mathbf{1}^\top$. Hence, for each node $k \in \mathcal{V}$, $\sum_{j=1}^n a_{kj} = \sum_{j=1}^n a_{jk}$, and, consequently, $\mathbf{1}^\top LB = 0$.

Let \tilde{x} be a fixed point of (6a), that is, $(I - A)\tilde{x} = LBv$. Since $A\mathbf{1} = \mathbf{1}$, it follows that $I = \text{diag}(\sum_j a_{1,j}, \dots, \sum_j a_{n,j})$, and $I - A = LB$. Then $LB(\tilde{x} - v) = 0$ for every fixed point \tilde{x} of (6a). Because A is a consensus matrix, 1 is a simple eigenvalue of A , and $A\mathbf{1} = \mathbf{1}$. Then $\text{Ker}(LB) = \text{Ker}(I - A) = \text{Im}(\mathbf{1})$. Notice that, by construction, $\eta = \frac{2\pi}{p} Bv$, $\tilde{\theta} = \frac{2\pi}{p} \tilde{x}$, and $B\mathbf{1} = 0$. Consequently, $(\tilde{\theta} - \eta) \in \text{Im}(\mathbf{1})$, and statement (ii) follows. \square

In Theorem 6.3 we assume that the measurements satisfy $\eta \in \text{Im}(B)$, or, equivalently, that the measurements are not affected by noise. While this assumption is justified by the fact that we only consider discretized measurements, in what follows we study the evolution of Algorithm 2 when the measurements are affected by noise. Let $e(t) = \eta - B\theta(t)$.

Theorem 6.4 (Convergence of Algorithm 2 with noisy measurements) *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a camera network,*

and let \mathcal{G}_d be its associated directed graph. Let $\eta \in \mathcal{O}_p^m$ be the vector of relative measurements, and let B be the incidence matrix of \mathcal{G}_d . If A achieves average consensus over \mathbb{F}_p , then

- (i) there exists a finite time $T < n$ such that the estimation error is constant, that is, $e(t) = e(t+1)$ for all $t \geq T$, and
- (ii) for all $t \geq T$, the estimation error satisfies

$$e(t) = \left(I - B \sum_{\tau=0}^{T-1} A^\tau L \right) \eta_{\text{orth}},$$

where η_{orth} is the orthogonal projection of η onto $\text{Im}(B)^\perp$, and L is as in (7).

PROOF: With the same notation as in the proof of Theorem 6.3, let $T < n$ be the number of iterations needed for convergence of the network matrix A , that is, $A^T = n^{p-2} \mathbf{1}\mathbf{1}^\top$. Notice that $BA^t = 0$ for all $t \geq T$, so that

$$e(t) = \eta - A^t x(0) - \sum_{\tau=0}^{t-1} A^\tau L \eta = \eta - \sum_{\tau=0}^{T-1} A^\tau L \eta,$$

for all $t \geq T$, and statement (i) follows. To show statement (ii), let $\eta = \eta_{\text{par}} + \eta_{\text{orth}}$, where $\eta_{\text{par}} \in \text{Im}(B)$, and $\eta_{\text{orth}} \in \text{Im}(B)^\perp$. From the linearity of (6a) and Theorem 6.3 we have

$$\eta_{\text{par}} = B \sum_{\tau=0}^{T-1} A^\tau L \eta_{\text{par}},$$

which concludes the proof. □

Theorem 6.4 characterizes the performance of Algorithm 2 when the measurements are affected by noise. Notice that the estimation error can be minimized by properly choosing the network matrix A . We now conclude with a numerical example.

Example 8 (Average computation and pose estimation via finite-field consensus) Consider a camera network with 4 cameras configured in a circle topology and network matrix

$$A = \begin{bmatrix} 0 & 4 & 2 & 0 \\ 1 & 1 & 0 & 4 \\ 0 & 0 & 2 & 4 \\ 0 & 1 & 2 & 3 \end{bmatrix} \in \mathbb{F}_5^{4 \times 4}.$$

It can be verified that A achieves average consensus over \mathbb{F}_5 in at most 3 iterations. In Fig. 5(a) we show that the network matrix A allows for the computation of the real-valued average of the agents initial states in 3 iterations (see Algorithm 1 and Section 4).

Let $A_k \in \mathbb{F}_5^{1024 \times 1024}$ be the network matrix generated from A as $A_k = A \otimes A \otimes A \otimes A \otimes A$. The sparsity pattern of A_k is reported in Fig. 5(b). In Fig. 6 we validate our distributed pose estimation algorithm (see Section 6.2). □

7 Conclusion and Future Work

In this paper we propose a distributed consensus algorithm for agents with limited memory, computation, and communication capabilities. Our approach is based on finite-fields, where agents states lie in a finite set, and operations are performed according to modular arithmetic. For our algorithm we identify necessary and sufficient convergence

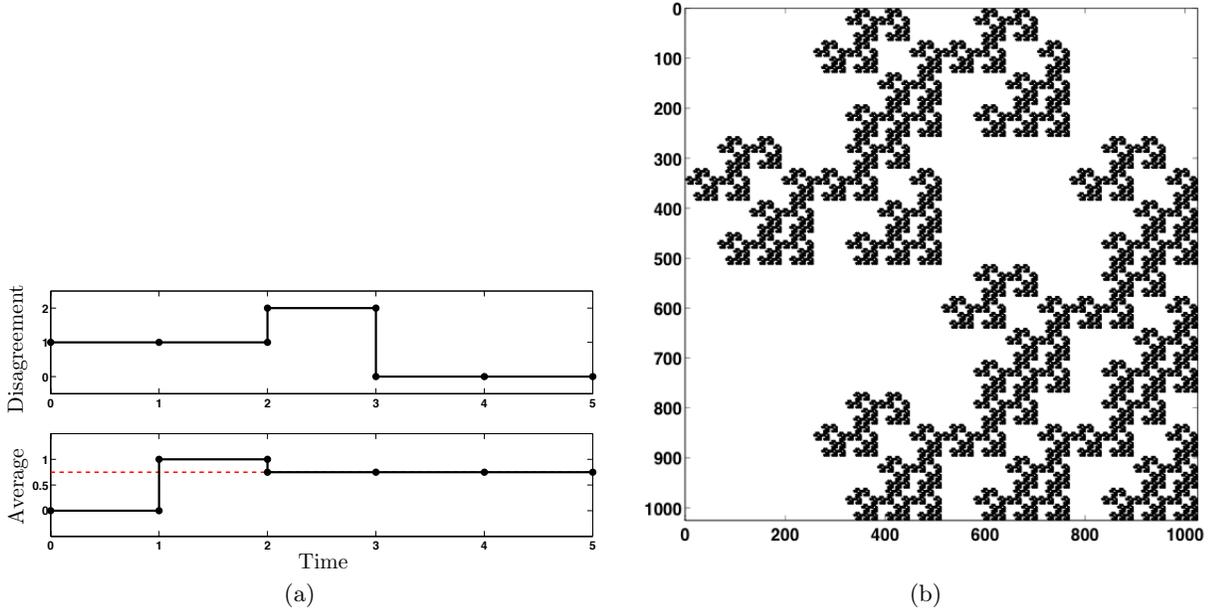


Fig. 5. For the network described by the matrix A in Example 8, let $0, 1, 1, 1$ be the agents initial states, respectively. Agents implement Algorithm 1. In Fig. 5(a) (above), we report the network disagreement as a function of time, where the disagreement at a given time equals the largest agents state minus the smallest agents state. Notice that the network achieves consensus in 3 iterations (see Section 4). In Fig. 5(a) (below), we report the average computed by the first agent (solid black) as a function of time (see Section 6.1). Notice that the first agent, and hence every agent in the network, computes the average of the initial states (dashed red) at the third iteration. Fig. 5(b) shows the sparsity pattern of the matrix $A_k \in \mathbb{F}_5^{1024 \times 1024}$ in Example 8. The network defined by A_k has 1024 nodes and 10^5 edges.

conditions, and we characterize the convergence time. Additionally, we discuss several network design methods, and we propose some application scenarios. Our work proposes a novel class of consensus dynamics, which are advantageous in several applications, and it complements the existing literature on real-valued consensus.

Through our analysis we show that finite-field consensus networks outperforms their real-valued counterpart in many aspects, including the convergence speed, robustness to communication errors, and agents requirements. These advantages come at the expenses of a more convoluted network design, which we identify as an interesting research direction. In particular, distributed network design algorithms, as well as gossip and asynchronous protocols should be investigated to broaden the applicability of finite-field consensus networks. Other theoretical research questions include characterizing the existence of finite-field consensus weights for a given interconnection structure and field characteristic, as well as the design of fastest finite-field consensus networks with fixed agents interconnection graph.

APPENDIX

Before proving Theorem 4.3, we recall the following fundamental results and facts in linear algebra.

Theorem 7.1 (Primary decomposition theorem [25]) *Let $A : \mathcal{V} \rightarrow \mathcal{V}$ be a linear operator on some vector space \mathcal{V} over some field \mathbb{F} , and let $p(s) = \prod_{i=1}^r p_i(s)$ be an annihilating polynomial for A with degree greater than 1, for some relatively prime polynomials p_1, \dots, p_r . Then*

- (i) $\mathcal{W}_i = \text{Ker}(p_i(A))$ is a A -invariant subspace for all $i \in \{1, \dots, r\}$,
- (ii) $\mathcal{V} = \mathcal{W}_1 \oplus \mathcal{W}_2 \oplus \dots \oplus \mathcal{W}_r$, where \oplus denote the direct sum operator, and
- (iii) if $\prod_{i=1}^r p_i(s) = P_A(s)$ and A_i is the restriction of A to \mathcal{W}_i , then p_i is the characteristic polynomial of A_i .

Recall that the order of a polynomial $g \in \mathbb{F}[s]$, denoted by $\text{ord}(g)$, is the smallest positive integer r such that $g(s)$ divides $s^r - 1$ over \mathbb{F} , that is, the smallest positive integer r such that there exists $q \in \mathbb{F}[s]$ satisfying $s^r - 1 = g(s)q(s)$.

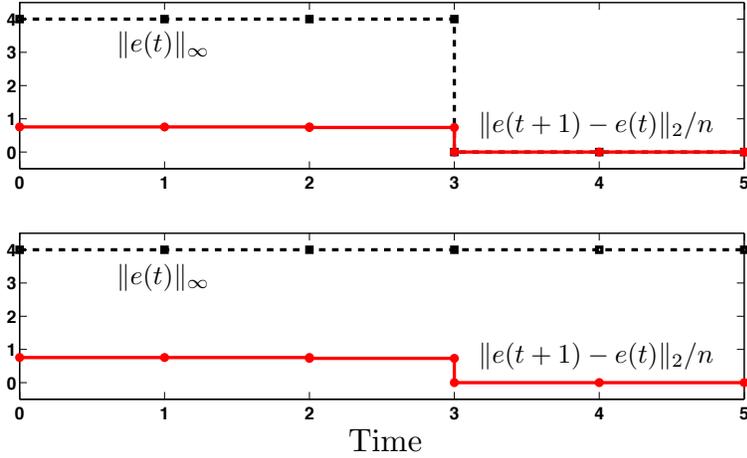


Fig. 6. For the camera network with 1024 cameras described by the matrix A_k in Example 8, this figure show the effectiveness of the our pose estimation algorithm in Algorithm 2. The figure above considers the noiseless case (Theorem 6.3). The dashed black line denotes the infinity norm of the estimation error, while solid red line corresponds to the (normalized) two norm of the difference between two consecutive estimation errors. We conclude that, in the absence of measurement noise, Algorithm 2 converges in 3 iterations to a configuration compatible with the relative measurements (Theorem 6.4). Notice that, since the measurements are affected by noise, the estimation error remains nonzero (dashed black). However, the algorithm converges to a configuration with constant estimation error. Indeed, the difference between two consecutive estimation errors is zero after 3 iterations (solid red).

Theorem 7.2 (Order of a polynomial over finite field [24]) *Let $g \in \mathbb{F}[s]$ be an irreducible polynomial satisfying $g(0) \neq 0$ and $\text{ord}(g) = e$. Consider $f = g^s$, and t is the smallest integer such that $p^t \geq s$, then $\text{ord}(f) = ep^t$.*

We are now ready to prove Theorem 4.3.

PROOF of Theorem 4.3: Let $P_A \in \mathbb{F}_p[s]$ be the characteristic polynomial of A , and notice that P_A can be written as

$$P_A(s) = \det(sI - A) = s^h \bar{P}(s), \quad (\text{A-1})$$

for some $h \in \mathbb{N}_{\geq 0}$, and $\bar{P}(s) \in \mathbb{F}_p[s]$ with $\bar{P}(0) \neq 0$.

(i) \implies (ii) Due to Lemma 3.2, we have $1 \in \sigma_p(A)$. Thus, we factorize P_A in irreducible polynomials as

$$P_A(s) = (s - 1)^k \prod_{j=1}^r Q_j(s)^{m_j},$$

where $k, m_j \in \mathbb{N}$ are given by the algebraic multiplicity of the corresponding eigenvalue.

We start by showing that $k = 1$. Assume by contradiction that $k > 1$. Let $\mathcal{W}_2 = \text{Ker}((I - A)^k)$, and let A_2 be the restriction of A to \mathcal{W}_2 . Recall from [27] that the cycle structure of the transition graph \mathcal{G}_2 of A_2 is

$$\text{Cycles}(\mathcal{G}_2) = C_1 + \sum_{i=1}^k (p^i - p^{i-1})C_1, \quad (\text{A-2})$$

where the sums of cycles is just the corresponding union graph, and C_1 denotes a unitary cycle, that is, a fixed point for A . From (A-2) it follows that, if $k > 1$, then the number of cycles in \mathcal{G}_2 is strictly greater than p . By Theorem 4.1 we conclude that $k = 1$.

We now show that $r = 0$. Assume by contradiction that $r > 0$. Let $\mathcal{W}_3 = \text{Ker}(Q_j(A)^{m_j})$, and let A_3 be the restriction of A to \mathcal{W}_3 . Then the cycle structure of the transition graph \mathcal{G}_3 of A_3 is

$$\text{Cycles}(\mathcal{G}_3) = C_1 + \sum_{i=1}^{m_j} \frac{p^{\deg(Q_j)^i} - p^{\deg(Q_j)^{i-1}}}{\ell_i} C_{\ell_i},$$

where $\ell_i = \text{ord}(Q_j^{n_j}) \geq \deg(Q_j) \geq 1$ from Theorem 7.2, and $\deg(\cdot)$ denotes the degree of a polynomial. Since the graph structure of A is given by the product of the graphs associated with the irreducible factors of its characteristic polynomial [27], the number of cycles is greater than p whenever either $k > 1$ or $r > 0$ (see Example 4 and Fig. 2).

(ii) \implies (i) Let $\mathcal{W}_1 = \text{Ker}(A - I) = \text{Im}(\mathbf{1})$ and recall from Theorem 7.1 that \mathcal{W}_1 is A -invariant. Let $V = [V_1 \mathbf{1}]$ be an invertible matrix, where the columns of V_1 are a basis for \mathcal{W}_1^\perp . Then we have

$$\tilde{A} = V^{-1}AV = \begin{bmatrix} A_{11} & 0 \\ A_{21} & 1 \end{bmatrix}.$$

Since the eigenvalues of a matrix are not affected by similarity transformations, the characteristic polynomial of the matrix A_{11} is s^{n-1} , so that A_{11} is nilpotent. It follows that every vector in \mathcal{W}_1^\perp converges to the origin in at most $n-1$ iterations, while vectors in \mathcal{W}_1 (consensus vectors) are fixed points for the matrix A . This concludes the proof. \square

References

- [1] N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1997.
- [2] F. Bullo, J. Cortés, and S. Martínez. *Distributed Control of Robotic Networks*. Applied Mathematics Series. Princeton University Press, 2009.
- [3] M. Mesbahi and M. Egerstedt, editors. *Graph Theoretic Methods in Multiagent Networks*. Princeton University Press, 2010.
- [4] F. Garin and L. Schenato. A survey on distributed estimation and control applications using linear consensus algorithms. In A. Bemporad, M. Heemels, and M. Johansson, editors, *Networked Control Systems*, LNCIS, pages 75–107. Springer, 2010.
- [5] W. Ren, R. W. Beard, and E. M. Atkins. Information consensus in multivehicle cooperative control: Collective group behavior through local interaction. *IEEE Control Systems Magazine*, 27(2):71–82, 2007.
- [6] L. Xiao, S. Boyd, and S. Lall. A scheme for robust distributed sensor fusion based on average consensus. In *Symposium on Information Processing of Sensor Networks*, pages 63–70, Los Angeles, CA, USA, April 2005.
- [7] D. P. Bertsekas and J. N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, 1997.
- [8] R. Lidl and H. Niederreiter. *Finite Fields*. Cambridge University Press, 1996.
- [9] Y. G. Sun, L. Wang, and G. Xie. Average consensus in networks of dynamic agents with switching topologies and multiple time-varying delays. *Systems & Control Letters*, 57(2):175–183, 2008.
- [10] T. C. Aysal, M. E. Yildiz, A. D. Sarwate, and A. Scaglione. Broadcast gossip algorithms for consensus. *IEEE Transactions on Signal Processing*, 57(7):2748–2761, 2009.
- [11] S. Kar and J. M. F. Moura. Distributed consensus algorithms in sensor networks with imperfect communication: Link failures and channel noise. *IEEE Transactions on Signal Processing*, 57(1):355–369, 2009.
- [12] L. Moreau. Stability of multiagent systems with time-dependent communication links. *IEEE Transactions on Automatic Control*, 50(2):169–182, 2005.
- [13] R. Carli and F. Bullo. Quantized coordination algorithms for rendezvous and deployment. *SIAM Journal on Control and Optimization*, 48(3):1251–1274, 2009.
- [14] A. Nedić, A. Olshevsky, A. Ozdaglar, and J. N. Tsitsiklis. On distributed averaging algorithms and quantization effects. *IEEE Transactions on Automatic Control*, 54(11):2506–2517, 2009.
- [15] R. Carli, F. Fagnani, P. Frasca, and S. Zampieri. Gossip consensus algorithms via quantized communication. *Automatica*, 46(1):70–80, 2010.
- [16] T. Li, M. Fu, L. Xie, and J. F. Zhang. Distributed consensus with limited communication data rate. *IEEE Transactions on Automatic Control*, 56(2):279–292, 2011.
- [17] J. Lavaei and R. M. Murray. Quantized consensus by means of gossip algorithm. *IEEE Transactions on Automatic Control*, 57(1):19–32, 2012.
- [18] A. Fagiolini, E. M. Visibelli, and A. Bicchi. Logical consensus for distributed network agreement. In *IEEE Conf. on Decision and Control*, pages 5250–5255, Cancún, México, December 2008.

- [19] A. Fagiolini, N. Dubbini, and A. Bicchi. Distributed consensus on set-valued information, 2011. Available at <http://arxiv.org/abs/1101.2275>.
- [20] A. Kashyap, T. Başar, and R. Srikant. Quantized consensus. *Automatica*, 43(7):1192–1203, 2007.
- [21] A. Olshevsky. Consensus with ternary messages, 2012. Available at <http://arxiv.org/abs/1212.5768>.
- [22] S. Sundaram and C. Hadjicostis. Structural controllability and observability of linear systems over finite fields with applications to multi-agent systems. *IEEE Transactions on Automatic Control*, 58(1):60–73, 2013.
- [23] R. Koetter and M. Médard. An algebraic approach to network coding. *IEEE/ACM Transactions on Networking*, 11(5):782–795, 2003.
- [24] B. Elspas. The theory of autonomous linear sequential networks. *IRE Transactions on Circuit Theory*, 6(1):45–60, 1959.
- [25] G. Shilov. *Linear Algebra*. New York: Dover Publications, 1977.
- [26] C. D. Godsil and G. F. Royle. *Algebraic Graph Theory*, volume 207 of *Graduate Texts in Mathematics*. Springer, 2001.
- [27] R. A. H. Toledo. Linear finite dynamical systems. *Communications in Algebra*, 33(9):2977–2989, 2005.
- [28] W. Ren and R. W. Beard. Consensus seeking in multi-agent systems under dynamically changing interaction topologies. *IEEE Transactions on Automatic Control*, 50(5):655–661, 2005.
- [29] J. C. Delvenne, R. Carli, and S. Zampieri. Optimal strategies in the average consensus problem. In *IEEE Conf. on Decision and Control*, New Orleans, USA, December 2007.
- [30] C. D. Meyer. *Matrix Analysis and Applied Linear Algebra*. SIAM, 2001.
- [31] P. Singla. On representations of general linear groups over principal ideal local rings of length two. *Journal of Algebra*, 324(9):2543–2563, 2010.
- [32] R. Olfati-Saber and R. M. Murray. Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on Automatic Control*, 49(9):1520–1533, 2004.
- [33] A. S. Fraenkel and Y. Yesha. Complexity of problems in games, graphs and algebraic equations. *Discrete Applied Mathematics*, 1(1-2):15–30, 1979.
- [34] M. R. Garey and D. S. Johnson. *Computers and Intractability*. Springer, 1979.
- [35] J. Ding, J. E. Gower, and D. Schmidt. *Multivariate Public Key Cryptosystems*. Springer, 2006.
- [36] L. Bettale, J. C. Faugère, and L. Perret. Hybrid approach for solving multivariate systems over finite fields. *Journal of Mathematical Cryptology*, 3(3):177–197, 2009.
- [37] N. Courtois, A. Klimov, J. Patarin, and A. Shamir. Efficient algorithms for solving overdefined systems of multivariate polynomial equations. In *Advances in Cryptology – EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 392–407. Springer, 2000.
- [38] J. Leskovec, D. Chakrabarti, J. Kleinberg, C. Faloutsos, and Z. Ghahramani. Kronecker graphs: An approach to modeling networks. *The Journal of Machine Learning Research*, 11:985–1042, 2010.
- [39] P. M. Weichsel. The Kronecker product of graphs. *Proceedings of the American Mathematical Society*, 13(1):47–52, 1962.
- [40] R. A. Horn and C. R. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, 1994.
- [41] S. Mac Lane. Modular fields. *The American Mathematical Monthly*, 47(5):259–274, 1940.
- [42] B. S. Y. Rao and H. F. Durrant-Whyte. A decentralized Bayesian algorithm for identification of tracked targets. *IEEE Transactions on Systems, Man & Cybernetics*, 23(6):1683–1698, 1993.
- [43] M. H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.
- [44] G. Piovan, I. Shames, B. Fidan, F. Bullo, and B. D. O. Anderson. On frame and orientation localization for relative sensing networks. *Automatica*, 49(1):206–213, 2013.
- [45] W. J. Russell, D. J. Klein, and J. P. Hespanha. Optimal estimation on the graph cycle space. *IEEE Transactions on Signal Processing*, 59(6):2834–2846, 2011.
- [46] D. Borra, E. Lovisari, R. Carli, F. Fagnani, and S. Zampieri. Autonomous calibration algorithms for networks of cameras. In *American Control Conference*, pages 5126–5131, Montréal, Canada, June 2012.