52nd IEEE Conference on Decision and Control
December 10-13, 2013. Florence, Italy

Finite-Field Consensus

Fabio Pasqualetti, Domenica Borra, and Francesco Bullo

***Abstract*— This work studies consensus networks over finite fields, where agents process and communicate values from the set of integers $\{0, \dots, p-1\}$, for some prime number p, and operations are performed modulo p. For consensus networks over finite fields we provide necessary and sufficient conditions on the network topology and weights to ensure convergence. For instance we show that, differently from the case of consensus networks over the field of real numbers, consensus networks over finite fields converge in finite time, and that properties of the agents interaction graph are not sufficient to ensure finite-field consensus. Finally, we discuss the application of finite-field consensus to distributed averaging in sensor networks.**

I. Introduction

Reaching agreement (consensus) is a fundamental task in distributed systems and networks [1]. Consensus algorithms have been used in several domains including robotics [2], estimation [3], and parallel computation [4]. In this work we focus on the consensus problem for networks of agents with limited resources and capabilities. In particular, we assume that agents process and communicate only values from a finite and pre-specified alphabet, and we 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 [5]. We focus on linear protocols, where agents update their state as a weighted combination of their neighbors states. Finite-field consensus finds applicability, for instance, in distributed averaging, load balancing, and pose estimation; it is easily implementable, and resilient to processing and communication noise [6].

Related work Consensus algorithms have been proposed for different network models, agents dynamics, and communication schemes [7]–[9], and detailed convergence conditions have been characterized [10]. While most of these approaches assume the possibility of processing and transmitting real values, we consider the case of finite communication bandwidth, and we show that certain topological conditions ensuring consensus over real values and with real-valued communications are not sufficient for finite-field consensus.

Consensus with *quantized* communication channels is studied, for instance, in [11]–[15]. In quantized consensus agents communicate quantized data, yet they perform computations over the field of real numbers. Thus, quantized consensus differs from finite-field consensus, where agents operate on a finite field.

This work was supported by NSF grants IIS-0904501 and CNS-1035917, and by ARO grant W911NF-11-1-0092. Fabio Pasqualetti is with the Mechanical Engineering Department, University of California, Riverside, fabiopas@engr.ucr.edu. Francesco Bullo is with the Center for Control, Dynamical Systems and Computation, University of California at Santa Barbara, bullo@engineering.ucsb.edu. Domenica Borra is with the Dipartimento di Scienze Matematiche, Politecnico di Torino, Italy, domenica.borra@polito.it.

Logical consensus is considered in [16] for the purpose of intruder and event detection. In logical consensus agents aim to coordinate their decisions via distributed computation of a function of a set of logical (boolean) events. Finite-field consensus differs from logical consensus in that (i) the agents state takes value in an arbitrary finite set, instead of being a binary variable, (ii) agents perform only two mathematical operations, namely modular addition and multiplication, and (iii) agents compute a non-boolean function of the states.

Consensus with integer communication and computation is studied in [17]. With respect to this work and to [18], we make use of *modular arithmetic*, instead of standard arithmetic, defining therefore a novel and complementary class of consensus networks. As discussed in [6], modular arithmetic is advantageous in several applications.

Finally, networks based on modular arithmetic are studied in [19], in the context of system controllability and observability, in [20], in the context of (linear) network coding, and in [21], in the context of finite dynamical systems.

Contributions The contributions of this paper are threefold.

First, we design distributed consensus networks based on finite fields and modular arithmetic (Section II). Consensus networks over finite fields are distributed, require limited memory, computation, and communication resources, and exhibit finite time convergence. Thus, finite-field consensus algorithms are suitable for capacity and memory constrained networks, and for time-constrained applications.

Second, we exhaustively characterize the convergence properties of consensus networks over finite fields (Section III). We provide necessary and sufficient constructive conditions on the agents interaction graph and weights to achieve finite-field 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 $s^{n-1}(s-1)$. Equivalently, consensus is achieved if and only if the transition graph of the network matrix contains exactly p cycles, where p is the field cardinality. 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 differs and complements the classic literature on real-valued consensus networks.

Third and finally, we discuss the case of finite-field average consensus, and we show how finite-field average consensus networks can be employed for the averaging problem in sensor networks over the field of real numbers.

II. 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 [5]. Let the directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$

describe the agents interaction graph, where $\mathcal{V} = \{1, \dots, n\}$, $i \in \mathcal{V}$ denotes the i-th agent, 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 allow each agent to process and manipulate values from the finite field $\mathbb{F}_{\mathrm{p}}$. We consider distributed laws where agent i updates its state $x_i \in \mathbb{F}_{\mathrm{p}}$ as a weighted combination of the states of its in-neighbors $\mathcal{N}_i^{\mathrm{in}} = \{w \in \mathcal{V} \ : \ (v, w) \in \mathcal{E}\}$. Define the network matrix $A \in \mathbb{F}_{\mathrm{p}}^{n \times n}$, $A = [a_{ij}]$, as the weighted adjacency matrix of $\mathcal{G}$, where $a_{ij} \in \mathbb{F}_{\mathrm{p}}$ is the weight associated with the edge (i, j), and $a_{ij} = 0$ if $(i, j) \notin \mathcal{E}$. Let $x : \mathbb{N}_{\geq 0} \to \mathbb{F}_{\mathrm{p}}^n$ be a map, where $x(t)$ is the vector of the agents states at time t. Then the network evolves as

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

where all operations are performed in the field $\mathbb{F}_{\mathrm{p}}$.

The *transition graph* associated with the iteration (1) over $\mathbb{F}_{\mathrm{p}}$ is defined as $\mathcal{G}_A = (\mathcal{V}_A, \mathcal{E}_A)$, where, $\mathcal{V}_A = \{v \ : \ v \in \mathbb{F}_{\mathrm{p}}^n\}$ and, for $v_i, v_j \in \mathcal{V}_A$, the edge $(v_i, v_j) \in \mathcal{E}_A$ if and only if $v_j = Av_i$. The transition graph contains p^n vertices and p^n edges, since each vertex has unit out-degree. Additionally, the transition graph is composed of disjoint weakly-connected subgraphs, each one containing a globally reachable node[1] and one cycle [23]. See Fig. 1 and Fig. 2.

The iteration (1) over a finite field achieves

(i) **asymptotic consensus**, if for all initial states $x(0) \in \mathbb{F}_{\mathrm{p}}$ it holds $\lim_{t \to \infty} x(t) = \alpha \mathbb{1}$, with $\alpha \in \mathbb{F}_{\mathrm{p}}$ and $\mathbb{1} = [1 \ \dots \ 1]^{\mathsf{T}}$;

(ii) **finite-time consensus**, if for all initial states $x(0) \in \mathbb{F}_{\mathrm{p}}$ there exists a finite time $T \in \mathbb{N}$ such that $x(T) = x(T + \tau) = \alpha \mathbb{1}$ for all $\tau \in \mathbb{N}$, with $\alpha \in \mathbb{F}_{\mathrm{p}}$ and $\mathbb{1} = [1 \ \dots \ 1]^{\mathsf{T}}$.

Consensus networks with real-valued weights and states have been extensively studied [1], [10], [24]. In this work we show that finite-field consensus networks differ from real-valued consensus networks, and particular care needs to be taken to ensure the desired properties over finite fields. Clearly, finite-time consensus implies asymptotic consensus. We next show that the converse is also true.

Theorem 2.1: ***(Asymptotic consensus implies finite-time consensus)*** The iteration (1) over the field $\mathbb{F}_{\mathrm{p}}$ 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, $\mathcal{C} = \{v \ : \ v \in \mathcal{V}_A, v = \alpha \mathbb{1}, \alpha \in \mathbb{F}_{\mathrm{p}}\}$. 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 and, equivalently, a finite number of iterations in (1) are sufficient to achieve consensus. ■

[1]A globally reachable node of a graph $\mathcal{G}$ is a vertex v to which there exists a directed path from every vertex in the graph, including v itself [22].

TABLE I

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

While consensus networks over finite fields either converge in finite time or they are not convergent (Theorem 2.1), 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 [25]. 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 $\mathbb{F}_{\mathrm{p}}$ is *nilpotent* if $A^n = 0$ and is *row-stochastic* if $A\mathbb{1} = \mathbb{1}$.

Lemma 2.2: ***(Finite-field consensus matrices)*** Consider the iteration (1) over the field $\mathbb{F}_{\mathrm{p}}$. If consensus is achieved, then A is either nilpotent or row-stochastic.

Proof: Since A achieves consensus, Theorem 2.1 implies that $A^t x(0) = A^{t+1} x(0) = \alpha \mathbb{1}$ for some $\alpha \in \mathbb{F}_{\mathrm{p}}$, for all $x(0)$, and for all $t \geq T$, $T \in \mathbb{N}$. Then $A\alpha\mathbb{1} = \alpha\mathbb{1}$, from which we conclude that either $A\mathbb{1} = \mathbb{1}$ (A is row-stochastic) or $\alpha = 0$ for all initial states $x(0)$ (A is nilpotent). ■

Following the above Lemma and as in the case of real-valued networks, we will only consider 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.

III. CONSENSUS NETWORKS OVER FINITE FIELDS

Convergence conditions for real-valued consensus are discussed, among others, in [1], [10], [24]. 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 finite-field consensus.

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

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

Notice that A_1, A_2, and A_3 are row-stochastic with fully connected interaction graph. It can be verified that only the network matrix A_1 achieves consensus over the field $\mathbb{F}_3$, 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 I. ■

As discussed in Example 1, graph properties of the network matrix are not sufficient to guarantee consensus for

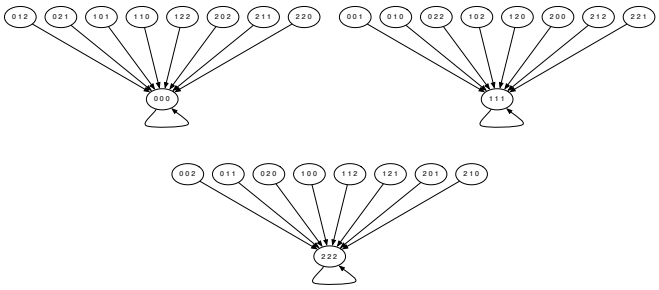

Fig. 1. Transition graph $\mathcal{G}_{A_1}$ associated with the matrix $A_1 \in \mathbb{F}_3^{3\times 3}$ in Example 1. Since $\mathcal{G}_{A_1}$ contains exactly 3 cycles corresponding to the self-loops around the consensus vertices, the network matrix A_1 achieves consensus (see Theorem 3.1).

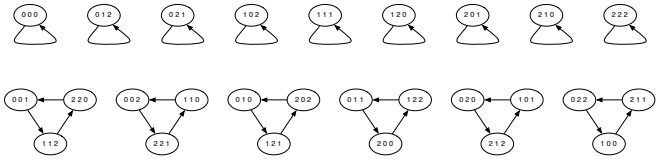

Fig. 2. Transition graph $\mathcal{G}_{A_3}$ associated with the matrix $A_3 \in \mathbb{F}_3^{3\times 3}$ in Example 1. Since $\mathcal{G}_{A_3}$ contains more than 3 cycles, the network matrix A_3 does not achieve consensus (see Theorem 3.1). The oscillatory state trajectory in Table I corresponds to the bottom right cycle in this figure.

iterations over a finite field. 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 transition graph of a finite-field network completely describes its dynamic behavior. The next theorem provides a necessary and sufficient condition for finite-field consensus based on the transition graph.

Theorem 3.1: ***(Transition graph of a consensus network)*** Consider the iteration (1) over the field $\mathbb{F}_\mathrm{p}$ 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:

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

Example 2: ***(Transition graph)*** 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, the matrix A_1 achieves consensus, while the matrix A_3 does not. ■

Proof of Theorem 3.1:

Notice that every state trajectory of the iteration (1) is in bijective correspondence with a path on its transition graph. *(i)* $\implies$ *(ii)* Since A is row-stochastic, it holds $A\mathbb{1} = \mathbb{1}$, and the transition graph of A contains p unit cycles corresponding to the vertices $\alpha\mathbb{1} \in \mathcal{V}_A$, with $\alpha \in \mathbb{F}_\mathrm{p}$. Suppose by contradiction that there exists an additional cycle C. Notice that the vertices $\alpha\mathbb{1}$, with $\alpha \in \mathbb{F}_\mathrm{p}$, cannot be contained in C since the out-degree of each vertex in the transition graph is one (the transition graph is determined by the linear map A). Thus, there exist state trajectories along C not converging to consensus, which contradicts the initial hypothesis.
(ii) $\implies$ *(i)* Suppose that transition graph $\mathcal{G}_A$ contains exactly p unit cycles located at the vertices $\alpha\mathbb{1} \in \mathcal{V}_A$, with $\alpha \in \mathbb{F}_\mathrm{p}$. Then, since each vertex in the transition graph has unit 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. ■

From condition (ii) in Theorem 3.1 and the fact that each vertex in the transition graph has unit 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 [23, Proposition 3.4] it can be shown that these disjoint subgraphs have the same graph topology.

A direct verification of the convergence condition in Theorem 3.1 may be prohibitive, due to the exponential growth of the size of the transition graph with the number of agents in the network (the transition graph contains p^n vertices and p^n edges, since each vertex has unit out-degree). In what follows we shall derive consensus conditions based on the network matrix instead of its transition graph. Consider the *inverse recursion*

$$\mathcal{S}_\alpha^{t+1} = A^{-1}(\mathcal{S}_\alpha^t) = \{x \in \mathbb{F}_\mathrm{p}^n \,:\, v = Ax, \forall v \in \mathcal{S}_\alpha^t\}, \quad (2)$$

where $\mathcal{S}_\alpha^0 = \{\alpha\mathbb{1}\}$ and $\mathcal{S}_\alpha^t \subset \mathbb{F}_\mathrm{p}^n$ for all times t. Notice that the inverse recursion defines a sequence of sets, and that the set $\mathcal{S}_\alpha^t$ contains the initial states converging to the consensus value $\alpha\mathbb{1}$ in at most t iterations. We say that the recursion (2) is convergent with limiting set $\mathcal{S}_\alpha$ 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 3.2: ***(Recursion sets of a consensus network)*** Consider the iteration (1) over the field $\mathbb{F}_\mathrm{p}$ with row-stochastic matrix A. The following statements are equivalent:

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

Example 3: ***(Inverse recursion)*** 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

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

Since $|\mathcal{S}_1| = 3^2$, the network matrix A_1 achieves consensus due to Theorem 3.2. Instead, for the network matrix $A_3 \in \mathbb{F}_3^{3\times 3}$ in Example 1, the inverse recursion yields $\mathcal{S}_1 = \{\mathbb{1}\}$, so that A_3 does not achieve consensus. ■

Proof of Theorem 3.2: 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 $\alpha\mathbb{1}$. Recall from [23, 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 these 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 unit cycles corresponding to the consensus vertices (see Theorem 3.1 and Fig. 1). By [23, Theorem 1], the above reasoning, and the fact that A achieves consensus, it follows that $|\mathcal{S}_0|+\cdots+|\mathcal{S}_{p-1}| = p^n$, and that $|\mathcal{S}_\alpha| = p^{n-1}$ for all $\alpha \in \mathbb{F}_\mathrm{p}$.

(ii) $\implies$ *(i)* Since $A\mathbb{1} = \mathbb{1}$, the transition graph contains p cycles of unit length located at the consensus vertices. Let the recursion (2) be convergent for some $\alpha \in \mathbb{F}_\mathrm{p}$. From [23, 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. ∎

From Theorem 3.2, it suffices to iterate the inverse recursion (2) for some $\alpha \in \mathbb{F}_\mathrm{p}$ to assess the convergence of the network (1). This method avoids the analysis of the transition graph. Our last and most explicit condition for finite-field consensus is based on the characteristic polynomial of the network matrix (over the finite field).

Theorem 3.3: ***(Characteristic polynomial of a consensus network)*** Consider the iteration (1) over the finite field $\mathbb{F}_\mathrm{p}$ 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)*** Consider the network matrices in Example 1 over the field $\mathbb{F}_3$. It can be verified that

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

As predicted by our previous analysis and by Theorem 3.3, only the network matrix A_1 achieves consensus. ∎

The proof of this theorem is postponsed to the Appendix. Theorem 3.3 is equivalently restated as follows: A achieves finite-field consensus if and only if $\sigma_\mathrm{p}(A) = \{1, 0, \ldots, 0\}$. 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; see [26] and [27, Theorem 3.5]. We conclude this section by characterizing the convergence value of a finite-field consensus network.

Theorem 3.4: ***(Finite-field consensus time and value)*** Consider the iteration (1) over the finite field $\mathbb{F}_\mathrm{p}$ with row-stochastic matrix A and with initial state $x(0) \in \mathbb{F}_\mathrm{p}^n$. Assume that the iteration (1) achieves consensus. Let $T < n$ denote the dimension of the largest Jordan block associated with the eigenvalue 0. Let $\pi \in \mathbb{F}_\mathrm{p}^n$ be the unique eigenvector satisfying $\pi A = \pi$ and $\pi\mathbb{1} = 1$. Then

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

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.[2]

Proof: Since A achieves consensus, we have $\sigma_\mathrm{p}(A) = \{1, 0, \ldots, 0\}$, and A admits a Jordan canonical form $J_A = V^{-1}AV$ over $\mathbb{F}_\mathrm{p}$. Moreover, the matrix A converges in $T < n$ iterations. The next part of the proof follows the reasoning in [28, Theorem 3]. Let the first column of V be $\mathbb{1}$, and let the matrix J_A have a unit 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 $A^T = V J_A^T V^{-1} = \mathbb{1}\pi$. Since $A\mathbb{1} = \mathbb{1}$, it follows that $\mathbb{1} = A^T\mathbb{1} = \mathbb{1}\pi\mathbb{1}$, and consequently $\pi\mathbb{1} = 1$. 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 $\mathbb{1}\pi x(0)$ needs to be zero, from which the statement follows. ∎

Note that Theorem 3.4 is not a direct consequence of the theory of non-negative matrices over the field of real numbers [26]. In fact, if regarded as a real-valued matrix, a finite-field consensus matrix is generally unstable.

IV. Application to Average Consensus

In this section we discuss the use of finite-field consensus for averaging (over real numbers) in sensor networks.

Consider a sensor network, and let $x_0 \in \mathbb{F}_\mathrm{p}^n$ denote the vector containing the agents initial states. Let $x_\mathbb{R} = \mathbb{1}^\mathsf{T}x_0/n \in \mathbb{R}$ be the average of the agents initial states over the field of real numbers. The average of the agents initial states over the field $\mathbb{F}_\mathrm{p}$ follows from Fermat's little theorem [29] as $x_\mathbb{F} = n^{p-2}\mathbb{1}^\mathsf{T}x_0 \in F_\mathrm{p}$, where we assume $n \neq kp$, with $k \in \mathbb{N}$, for the inverse of n over $\mathbb{F}_\mathrm{p}$ 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 to recover the average $x_\mathbb{R}$ from the knowledge of $x_\mathbb{F}$ and the total number of agents n. We say that the iteration (1) over the field $\mathbb{F}_\mathrm{p}$ achieves average consensus if it achieves consensus, and the consensus value is $n^{p-2}\mathbb{1}^\mathsf{T}x_0$ for every initial state x_0.

Theorem 4.1: ***(Finite-field average consensus)*** Consider the iteration (1) over the field $\mathbb{F}_\mathrm{p}$ 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 $\mathbb{1}^\mathsf{T}A = \mathbb{1}^\mathsf{T}$.

Example 5: ***(A finite-field average consensus network)*** Consider the network matrix

[2] A root node of a graph $\mathcal{G}$ is a vertex v from which there exists a directed path to every vertex in the graph, including v itself [22].

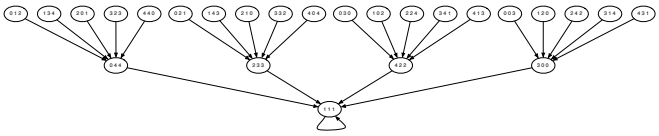

Fig. 3. A subgraph of the transition graph associated with the network matrix A in Example 5. Notice that the sum of the initial states is maintained, and thus average consensus is achieved.

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

over the field $\mathbb{F}_5$. It can be verified that $A\mathbb{1} = \mathbb{1}$, $\mathbb{1}^\mathsf{T} A = \mathbb{1}^\mathsf{T}$, and $P_A = s^2(s-1)$. By Theorem 4.1 the network matrix A achieves average consensus over $\mathbb{F}_5$. In Fig. 3 we show a subgraph of the transition graph associated with A. ■

Proof of Theorem 4.1:
(i) $\implies$ *(ii)* Since the iteration achieves consensus, it follows from Theorem 3.3 that $P_A(s) = s^{n-1}(s-1)$, and from Theorem 3.4 that $A^n = \mathbb{1}\pi$, where π satisfies $\pi A = \pi$. Because A achieves average consensus, it needs to be $\mathbb{1}\pi = n^{p-2}\mathbb{1}\mathbb{1}^\mathsf{T}$. Then $\pi = n^{p-2}\mathbb{1}^\mathsf{T}$, and $\mathbb{1}^\mathsf{T} A = \mathbb{1}^\mathsf{T}$.
(ii) $\implies$ *(i)* Because A is row-stochastic and $P_A(s) = s^{n-1}(s-1)$, the network achieves consensus due to Theorem 3.3. Notice that $\mathbb{1}^\mathsf{T} A = \mathbb{1}^\mathsf{T}$ implies that $\mathbb{1}^\mathsf{T} x(t) = \mathbb{1}^\mathsf{T} x(0)$ at all times t. Let α be the consensus value, and notice that $\mathbb{1}^\mathsf{T}\alpha\mathbb{1} = n\alpha = \mathbb{1}^\mathsf{T} x(0)$. Finally, $\alpha = n^{p-2}\mathbb{1}^\mathsf{T} x(0)$, and the network achieves average consensus. ■

Theorem 4.1 provides a necessary and sufficient condition for a network with $n \neq kp$ agents to achieve average consensus over $\mathbb{F}_\mathrm{p}$. The condition $n \neq kp$ is actually necessary for average consensus. In fact, if $n = kp$ for some $k \in \mathbb{N}$, then there exists no network matrix satisfying all conditions in Theorem 4.1 and, therefore, average consensus cannot be achieved. To see this, let $x(0)$ be the network initial state, with $\mathbb{1}^\mathsf{T} x(0) \neq 0$, and assume by contradiction that α is the corresponding consensus value. Since $\mathbb{1}^\mathsf{T} x(t) = \mathbb{1}^\mathsf{T} x(0)$ at all times t, it needs to be $n\alpha = \mathbb{1}^\mathsf{T} x(0)$. Then $0 = n\alpha = \mathbb{1}^\mathsf{T} x(0) \neq 0$, since kp and 0 are the same element in $\mathbb{F}_\mathrm{p}$.

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 $x_1 = [2\ 2\ 2]^\mathsf{T}$ and $x_2 = [0\ 0\ 1]^\mathsf{T}$. Over the field of real numbers we have $x_{1,\mathbb{R}} = \mathbb{1}^\mathsf{T} x_1/n = 2$ and $x_{2,\mathbb{R}} = \mathbb{1}^\mathsf{T} x_2/n = 1/3$. Over the field $\mathbb{F}_5$, instead, $x_{1,\mathbb{F}} = n^{p-2}\mathbb{1}^\mathsf{T} x_1 = 2$ and $x_{2,\mathbb{F}} = n^{p-2}\mathbb{1}^\mathsf{T} x_2 = 2$. 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. The next theorem contains a sufficient condition to recover the real-valued average from its finite-field counterpart.

Theorem 4.2: ***(Average computation)*** Let $x_0 \in \mathbb{F}_\mathrm{p}^n$, let $x_\mathbb{R} = \mathbb{1}^\mathsf{T} x_0/n \in \mathbb{R}$, and let $x_\mathbb{F} = n^{p-2}\mathbb{1}^\mathsf{T} x_0$. If the field characteristic satisfies $n\|x_0\|_\infty \leq p$, then

$$n\, x_\mathbb{R} = \mathrm{mod}(n\, x_\mathbb{F}, p),$$

where $\mathrm{mod}(\cdot)$ denotes the modulus operation.

Proof: The statement follows from the relation

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

where we have used $\mathrm{mod}(n^{p-1}, p) = 1$ and $n\|x_0\|_\infty \leq p$. ■

A finite-time averaging algorithm is readily derived from Theorem 4.1 and Theorem 4.2. We conclude by noticing that the condition $n\|x_0\|_\infty \leq p$ in Theorem 4.2 is not restrictive. In fact, the field characteristic p is a design parameter and it can be chosen to satisfy the above condition if the network cardinality and a bound on the agents initial state are known.

V. Conclusion and Future Work

In this work we study consensus networks over finite fields, where agents process and communicate values from the set $\{0, \ldots, p-1\}$, for some prime number p, and operations are performed modulo p. For finite-field consensus we identify necessary and sufficient convergence conditions, and we characterize several properties including the convergence time. Finite-field consensus is a novel consensus mechanism, which is advantageous for estimation and coordination problems in capacity and memory constrained networks.

APPENDIX

Before proving Theorem 3.3, we recall the following fundamental results and facts in linear algebra.

Theorem 5.1: ***(Primary decomposition theorem [30])*** Let $A : \mathcal{V} \to \mathcal{V}$ 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 $p_1, \ldots, p_r$. Then

(i) $\mathcal{W}_i = \mathrm{Ker}(p_i(A))$ is a A-invariant subspace for all $i \in \{1, \ldots, r\}$,[3]
(ii) $\mathcal{V} = \mathcal{W}_1 \oplus \mathcal{W}_2 \oplus \cdots \oplus \mathcal{W}_r$, 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 $\mathrm{ord}(g)$, 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 $s^r - 1 = g(s)q(s)$.

Theorem 5.2: ***(Order of a polynomial over a finite field [21])*** Let $g \in \mathbb{F}[s]$ be an irreducible polynomial satisfying $g(0) \neq 0$ and $\mathrm{ord}(g) = e$. Consider $f = g^s$, and t is the smallest integer such that $p^t \geq s$, then $\mathrm{ord}(f) = ep^t$.

We are now ready to prove Theorem 3.3.

Proof of Theorem 3.3: Let $P_A \in \mathbb{F}_\mathrm{p}[s]$ be the characteristic polynomial of A, and notice that P_A can be written as

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

[3] $\mathrm{Ker}(A) = \{x \in \mathbb{F}_\mathrm{p}^m : Ax = 0\}$ and $\mathrm{Im}(A) = \{y \in \mathbb{F}_\mathrm{p}^n : y = Ax, x \in \mathbb{F}_\mathrm{p}^m\}$ are the null space and the range space of the matrix $A \in \mathbb{F}_\mathrm{p}^{n \times m}$, respectively.

for some $h \in \mathbb{N}_{\geq 0}$, and $\bar{P}(s) \in \mathbb{F}_{\mathrm{p}}[s]$ with $\bar{P}(0) \neq 0$.
(i) $\implies$ *(ii)* Since A is row-stochastic, we have $1 \in \sigma_{\mathrm{p}}(A)$. 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 $\mathcal{W}_2 = \mathrm{Ker}((I-A)^k)$, and let A_2 be the restriction of A to $\mathcal{W}_2$. Recall from [23] that the cycle structure of the transition graph $\mathcal{G}_2$ of A_2 is

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

where the sum of cycles is simply the corresponding union graph, and C_1 denotes a unit 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 3.1 we conclude that $k = 1$.

We now show that $r = 0$. Assume by contradiction that $r > 0$. Let $\mathcal{W}_3 = \mathrm{Ker}(Q_j(A)^{m_j})$, 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

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

where $\ell_i = \mathrm{ord}(Q_j^{n_j}) \geq \deg(Q_j) \geq 1$ from Theorem 5.2, 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 [23], 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 $\mathcal{W}_1 = \mathrm{Ker}(A-I) = \mathrm{Im}(\mathbb{1})$ and recall from Theorem 5.1 that $\mathcal{W}_1$ is A-invariant. Let $V = [V_1 \, \mathbb{1}]$ be an invertible matrix, where the columns of V_1 are a basis for $\mathcal{W}_1^{\perp}$. Then we have

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

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. ■

REFERENCES

[1] F. Garin and L. Schenato, "A survey on distributed estimation and control applications using linear consensus algorithms," in *Networked Control Systems*, ser. LNCIS, A. Bemporad, M. Heemels, and M. Johansson, Eds. Springer, 2010, pp. 75–107.

[2] 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*, vol. 27, no. 2, pp. 71–82, 2007.

[3] 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*, Los Angeles, CA, USA, Apr. 2005, pp. 63–70.

[4] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, 1997.

[5] R. Lidl and H. Niederreiter, *Finite Fields*. Cambridge University Press, 1996.

[6] F. Pasqualetti, D. Borra, and F. Bullo, "Consensus networks over finite fields," *Automatica*, Jan. 2013, to appear.

[7] 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*, vol. 57, no. 2, pp. 175–183, 2008.

[8] T. C. Aysal, M. E. Yildiz, A. D. Sarwate, and A. Scaglione, "Broadcast gossip algorithms for consensus," *IEEE Transactions on Signal Processing*, vol. 57, no. 7, pp. 2748–2761, 2009.

[9] 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*, vol. 57, no. 1, pp. 355–369, 2009.

[10] L. Moreau, "Stability of multiagent systems with time-dependent communication links," *IEEE Transactions on Automatic Control*, vol. 50, no. 2, pp. 169–182, 2005.

[11] R. Carli and F. Bullo, "Quantized coordination algorithms for rendezvous and deployment," *SIAM Journal on Control and Optimization*, vol. 48, no. 3, pp. 1251–1274, 2009.

[12] A. Nedić, A. Olshevsky, A. Ozdaglar, and J. N. Tsitsiklis, "On distributed averaging algorithms and quantization effects," *IEEE Transactions on Automatic Control*, vol. 54, no. 11, pp. 2506–2517, 2009.

[13] R. Carli, F. Fagnani, P. Frasca, and S. Zampieri, "Gossip consensus algorithms via quantized communication," *Automatica*, vol. 46, no. 1, pp. 70–80, 2010.

[14] T. Li, M. Fu, L. Xie, and J. F. Zhang, "Distributed consensus with limited communication data rate," *IEEE Transactions on Automatic Control*, vol. 56, no. 2, pp. 279–292, 2011.

[15] J. Lavaei and R. M. Murray, "Quantized consensus by means of gossip algorithm," *IEEE Transactions on Automatic Control*, vol. 57, no. 1, pp. 19–32, 2012.

[16] A. Fagiolini, E. M. Visibelli, and A. Bicchi, "Logical consensus for distributed network agreement," in *IEEE Conf. on Decision and Control*, Cancún, México, Dec. 2008, pp. 5250–5255.

[17] A. Kashyap, T. Başar, and R. Srikant, "Quantized consensus," *Automatica*, vol. 43, no. 7, pp. 1192–1203, 2007.

[18] A. Olshevsky, "Consensus with ternary messages," 2012, available at `http://arxiv.org/abs/1212.5768`.

[19] 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*, vol. 58, no. 1, pp. 60–73, 2013.

[20] R. Koetter and M. Médard, "An algebraic approach to network coding," *IEEE/ACM Transactions on Networking*, vol. 11, no. 5, pp. 782–795, 2003.

[21] B. Elspas, "The theory of autonomous linear sequential networks," *IRE Transactions on Circuit Theory*, vol. 6, no. 1, pp. 45–60, 1959.

[22] C. D. Godsil and G. F. Royle, *Algebraic Graph Theory*, ser. Graduate Texts in Mathematics. Springer, 2001, vol. 207.

[23] R. A. H. Toledo, "Linear finite dynamical systems," *Communications in Algebra*, vol. 33, no. 9, pp. 2977–2989, 2005.

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

[25] 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, Dec. 2007.

[26] C. D. Meyer, *Matrix Analysis and Applied Linear Algebra*. SIAM, 2001.

[27] P. Singla, "On representations of general linear groups over principal ideal local rings of length two," *Journal of Algebra*, vol. 324, no. 9, pp. 2543–2563, 2010.

[28] 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.

[29] S. M. Lane, "Modular fields," *The American Mathematical Monthly*, vol. 47, no. 5, pp. 259–274, 1940.

[30] G. Shilov, *Linear Algebra*. New York: Dover Publications, 1977.