Automatica 50 (2014) 349–358

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Consensus networks over finite fields ☆

Fabio Pasqualetti [a,1], Domenica Borra [b], Francesco Bullo [c]

[a] Department of Mechanical Engineering, University of California, Riverside, USA
[b] Dipartimento di Scienze Matematiche, Politecnico di Torino, Italy
[c] Center for Control, Dynamical Systems, and Computation, University of California, Santa Barbara, USA

ARTICLE INFO

Article history:
Received 16 January 2013
Received in revised form
19 May 2013
Accepted 7 October 2013
Available online 11 December 2013

Keywords:
Multi-agent system
Consensus network
Finite field
Distributed algorithm
Graph theory

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 $\{0, \ldots, p-1\}$, 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.

Published by Elsevier Ltd.

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 (Bullo, Cortés, & Martínez, 2009). An important example is the *consensus* problem, where members of a network aim to agree upon a parameter of interest via distributed computation (Garin & Schenato, 2010). Consensus algorithms have applicability in many domains, including robotics (Ren, Beard, & Atkins, 2007), estimation (Xiao, Boyd, & Lall, 2005), and parallel computation (Bertsekas & Tsitsiklis, 1997).

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. 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 (Lidl & Niederreiter, 1996). 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 is applicable 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 developed for different network models, 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 with time-varying topologies (Sun, Wang, & Xie, 2008), gossip and asynchronous communication (Aysal, Yildiz, Sarwate, & Scaglione, 2009), and communication errors and link failures (Kar & Moura, 2009). It has been shown that, under mild connectivity assumptions on the interaction graph, a simple linear iteration suffices to ensure consensus (Moreau, 2005). 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 and memory constraints. As we show, topological conditions ensuring consensus with real values and unlimited bandwidth are not sufficient for consensus over a finite field.

☆ This work was supported by NSF grants IIS-0904501 and CNS-1035917, and by ARO grant W911NF-11-1-0092. The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Tamas Keviczky under the direction of Editor Frank Allgöwer.

E-mail addresses: fabiopas@engr.ucr.edu (F. Pasqualetti), domenica.borra@polito.it (D. Borra), bullo@engineering.ucsb.edu (F. Bullo).

[1] Tel.: +1 951 827 2327.

0005-1098/$ – see front matter. Published by Elsevier Ltd.
http://dx.doi.org/10.1016/j.automatica.2013.11.011

A relevant body of the 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 (Li, Fu, Xie, & Zhang, 2011; Nedić, Olshevsky, Ozdaglar, & Tsitsiklis, 2009). In these works, although the exchanged data is quantized, agents perform real-valued computations. Thus, the consensus problem with quantized information is related to our problem, yet fundamentally different because we allow agents to operate only on a finite field.

Logical consensus has been studied in Fagiolini and Bicchi (2013) and Fagiolini, Visibelli, and Bicchi (2008) for the purpose of intruder and event detection. In logical consensus agents aim 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 Fagiolini and Bicchi (2013) and Fagiolini et al. (2008) is on the design of a synthesis technique for logical consensus systems. The main differences between Fagiolini et al. (2008) 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}, as opposed to only modular addition and multiplication. Third and finally, in logical consensus agents aim to agree upon a logical expression or compact sets, while finite-field consensus algorithms may be used to compute a (non-boolean) function, such as the average of the initial states.

A distributed consensus algorithm with integer communication and computation is proposed in Kashyap, Başar, and Srikant (2007). With respect to this work 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 Sundaram and Hadjicostis (2013), in the context of system controllability and observability, in Koetter and Médard (2003), in the context of (linear) network coding, and in Elspas (1959), in the more general context of finite dynamical systems.

Contributions: The contributions of this paper are fourfold. First, we propose the use of finite fields to design consensus algorithms for networks of cooperative agents (Section 3). Finite-field consensus networks are distributed, require limited, in fact finite, memory, computation, and communication resources, and converge in finite time. Thus, finite-field consensus algorithms are suitable for capacity and memory constrained networks, and for applications subject to time constraints.

Second, we characterize convergence of consensus networks over finite fields (Section 4). 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.

Third, we propose systematic methods to design consensus networks over finite fields (Section 5). In particular, we derive a general design method, 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 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 (Section 6). In the averaging problem agents aim to determine the average (over the field of real numbers) of their initial values. We show how, under a reasonable set of assumptions, the averaging problem can be solved distributively and in finite time via a finite-field average consensus. In the pose estimation problem agents aim to estimate their pose based on relative measurements. We derive a distributed pose estimation algorithm based on finite-field average consensus, and we characterize its performance.

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 Godsil and Royle (2001), Lidl and Niederreiter (1996) and Shilov (1977) for a comprehensive treatment of these subjects.

A *field* $\mathbb{F}$ is a set of elements with *addition* and *multiplication* operations, and satisfies the following axioms:

(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 elements $0, 1 \in \mathbb{F}$ such that $a+0 = a$ and $a \cdot 1 = a$;
(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 $\mathbb{F}_{\mathrm{p}}$ with characteristic p a prime number. The field $\mathbb{F}_{\mathrm{p}}$ 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.

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

$$\mathrm{Im}(A) := \{y \in \mathbb{F}_{\mathrm{p}}^n : y = Ax, x \in \mathbb{F}_{\mathrm{p}}^m\},$$

$$\mathrm{Ker}(A) := \{x \in \mathbb{F}_{\mathrm{p}}^m : Ax = 0\},$$

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

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

Let $\mathbb{F}_{\mathrm{p}}[s]$ denote the set of polynomials with coefficients in $\mathbb{F}_{\mathrm{p}}$, and let $P_A \in \mathbb{F}_{\mathrm{p}}[s]$ denote the characteristic polynomial of $A \in \mathbb{F}_{\mathrm{p}}^{n\times n}$ over $\mathbb{F}_{\mathrm{p}}$.[2] Let $\sigma_{\mathrm{p}}(A)$ denote the set of eigenvalues of A, that is, the

[2] Since the characteristic polynomial $\bar{P}_A(s) \in \mathbb{R}[s]$ has only integer coefficients for any $A \in \mathbb{F}_{\mathrm{p}}^{n\times n}$, the characteristic polynomial $P_A(s) \in \mathbb{F}_{\mathrm{p}}[s]$ is $\sum_{i=0}^{n} \bmod (\bar{c}_i, p) s^i$, where $\bmod(\cdot)$ is the modulus function, and $\bar{c}_i$ is the ith coefficient of $\bar{P}_A(s)$.