Automatica 48 (2012) 2344–2351

Contents lists available at SciVerse ScienceDirect

Automatica

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

Brief paper

Convergence time analysis of quantized gossip consensus on digraphs[☆]

Kai Cai[a,1], Hideaki Ishii[b,2]

[a] Department of Electrical and Computer Engineering, University of Toronto, 10 King's College Rd, Toronto, M5S 3G4, Canada

[b] Department of Computational Intelligence and Systems Science, Tokyo Institute of Technology, 4259-J2-54, Nagatsuta-cho, Midori-ku, Yokohama 226-8502, Japan

ARTICLE INFO

Article history:
Received 20 October 2010
Received in revised form
17 November 2011
Accepted 17 April 2012
Available online 11 July 2012

Keywords:
Quantized consensus
Average consensus
Convergence time
Randomized gossip algorithms
Directed graphs
Markov chains

ABSTRACT

We have recently proposed quantized gossip algorithms which solve the consensus and averaging problems on directed graphs with the least restrictive connectivity requirements. In this paper we study the convergence time of these algorithms. To this end, we investigate the shrinking time of the smallest interval that contains all states for the consensus algorithm, and the decay time of a suitable Lyapunov function for the averaging algorithm. The investigation leads us to characterize the convergence time by the hitting time in certain special Markov chains. We simplify the structures of state transition by considering the special case of complete graphs, where every edge can be activated with an equal probability, and derive polynomial upper bounds on convergence time.

1. Introduction

Inspired by aggregate behavior of animal groups and motion coordination of robotic networks, the *consensus* problem has been extensively studied in the recent literature of systems control (e.g., Jadbabaie, Lin, and Morse (2003), Olfati-Saber and Murray (2004) and Ren and Beard (2005)). The objective of consensus is to have a population of nodes, each possessing an initial state, agree eventually on *some* common value through only local information exchange. This problem is also related intimately to oscillator synchronization and load balancing. The *averaging* problem is of a special form, with the goal to decentrally compute the *average* of all initial states at every node.

We have recently proposed in Cai and Ishii (2011a,b) randomized gossip algorithms which solve the consensus and averaging problems on directed graphs (or digraphs), under a quantization constraint that each node has an integer-valued state. In particular, our derived connectivity condition ensuring average consensus is weaker than those in the literature (Olfati-Saber & Murray, 2004), in the sense that it does not postulate balanced topologies. Here the main difficulty is that the state sum of nodes cannot be preserved during algorithm iterations. This scenario was previously considered in Aysal, Coates, and Rabbat (2008); Aysal, Yildiz, Sarwate, and Scaglione (2009), where averaging is guaranteed in expectation but there is in general an error in mean square and with probability one. By contrast, we overcome this difficulty by augmenting the so-called "surplus" variables for individual nodes so as to maintain local records of state updates, thereby ensuring average consensus almost surely.

In this paper and its conference precursor (Cai & Ishii, 2010), we investigate the performance of our proposed algorithms by providing upper bounds on the *mean convergence time*. The state transition structures resulting from these algorithms turn out to be rather complicated. Hence in our analysis on convergence time, we focus on the special case of complete graphs whose topology is undirected. The analysis is still challenging, but we will also discuss that the general approach can be useful for other graph topologies. First, for the consensus algorithm, we find that the mean convergence time is $O(n^2)$. To derive this bound, we view reaching consensus as the smallest interval containing all states shrinking its length to zero. This perspective leads us to characterizing convergence time by the *hitting* time in a certain Markov chain, which yields the polynomial bound. Second, we obtain that the mean convergence time of the averaging algorithm is $O(n^3)$. As the original algorithm in Cai and Ishii (2011a,b) is found to induce complex state transition structures, we have suitably revised it to manage the complexity. In particular, a bidirectional communication protocol is introduced. For the modified algorithm, a Lyapunov function is proposed which measures the distance from average consensus. We then bound convergence time by way

☆ This work was supported in part by the Ministry of Education, Culture, Sports, Science and Technology in Japan under Grant-in-Aid for Scientific Research, No. 21760323. The material in this paper was partially presented at the 49th IEEE Conference on Decision and Control (CDC 2010), December 15–17, 2010, Atlanta, Georgia, USA. This paper was recommended for publication in revised form by Associate Editor Valery Ugrinovskii under the direction of Editor Ian R. Petersen.

E-mail addresses: kai.cai@scg.utoronto.ca (K. Cai), ishii@dis.titech.ac.jp (H. Ishii).

[1] Tel.: +1 416 978 6299; fax: +1 416 978 6299.

[2] Tel.: +81 45 924 5371.

0005-1098/$ – see front matter © 2012 Elsevier Ltd. All rights reserved.
doi:10.1016/j.automatica.2012.06.048

of bounding the number of iterations required to decrease the Lyapunov function; the latter is again characterized by the hitting time in a special Markov chain.

Our work is related to Lavaei and Murray (2012), Kashyap, Başar, and Srikant (2007) and Zhu and Martínez (2008), who also tackle the convergence time of gossip algorithms with quantized states. In Kashyap et al. (2007), polynomial bounds on convergence time are obtained for fully connected and linear networks. The work (Zhu & Martínez, 2008) generalizes these bounds to arbitrarily connected networks (fixed or switching). Also, bounds for arbitrarily connected networks are provided in Lavaei and Murray (2012); these bounds are in terms of graph topology. In these references, a common feature is that the graphs are undirected. To bound the convergence time, a frequently employed approach is to bound the decay time of a Lyapunov function (Kashyap et al., 2007; Nedic, Olshevsky, Ozdaglar, & Tsitsiklis, 2009). In particular, Nedic et al. (2009) derives tight polynomial bounds on the convergence time of synchronized averaging algorithms, with either real or quantized states. In addition, Carli, Fagnani, Frasca, and Zampieri (2010) investigates a variety of quantization effects on averaging algorithms, and demonstrates favorable convergence properties by simulations. Our work adopts the Lyapunov method, as in Kashyap et al. (2007) and Nedic et al. (2009); the common function used in these papers turns out, however, not to be a valid Lyapunov function for our averaging algorithm. This is due again to that the state sum does not remain invariant, and the augmented surplus evolution must also be taken into account. According to these features, we establish an appropriate Lyapunov function, and prove that bounding its decay time can be reduced to finding the hitting time in a certain Markov chain.

1.1. Setup and organization

Consider a digraph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{1, \ldots, n\}$ is the node set, and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ the edge set. Each directed edge (j, i) in $\mathcal{E}$, pointing from j to i, denotes that node j communicates to node i (namely, the information flow is from j to i). Communication among the nodes is by means of *gossip*: At each time instant, exactly one edge $(j, i) \in \mathcal{E}$ is activated independently from all earlier instants and with a time-invariant positive probability $p_{ji} \in (0, 1)$ such that $\sum_{(j,i)\in\mathcal{E}} p_{ji} = 1$.

To model the quantization effect in information flow, we assume that at time $k \in \mathbb{Z}_+$ (nonnegative integers), each node has an integer-valued state $x_i(k) \in \mathbb{Z}$, $i \in \mathcal{V}$; the aggregate state is denoted by $x(k) = [x_1(k) \cdots x_n(k)]^T \in \mathbb{Z}^n$. Let

$$\mathcal{X} := \{x : m \leq x_i \leq M,\ i \in \mathcal{V}\}, \tag{1}$$

for some (finite) constants m, M. Suppose throughout the paper that the initial state satisfies $x(0) \in \mathcal{X}$. Also, let $\mathbf{1} = [1 \cdots 1]^T$ be the vector of all ones.

For the convergence time analysis, we will impose the following two assumptions on the graph topology and the probability distribution of activating edges. Let $|\cdot|$ denote the cardinality of a set.

Assumption 1. The digraph $\mathcal{G}$ is *complete* (i.e., every node is connected to every other node by a directed edge). It follows that there are $|\mathcal{E}| = n(n-1)$ edges.

Assumption 2. The probability distribution on edge activation is *uniform*; namely, each edge can be activated with the same probability $p := 1/|\mathcal{E}|$.

The rest of this paper is organized as follows. In Section 2, we address the convergence time analysis for the consensus algorithm. Then in Sections 3 and 4, we derive an upper bound for the convergence time of the averaging algorithm. Further, we compare convergence rates through a numerical example in Section 5, and finally, we state our conclusions in Section 6.

2. Convergence time of consensus algorithm

2.1. Algorithm and problem formulation

First we recall the quantized consensus (**QC**) algorithm from Cai and Ishii (2011b). Suppose that the edge $(j, i) \in \mathcal{E}$ is randomly activated at time k. Along the edge node j sends to i its state information, $x_j(k)$, but does not perform any update, i.e., $x_j(k+1) = x_j(k)$. On the other hand, node i receives j's state $x_j(k)$ and updates its own as follows:

(R1) If $x_i(k) = x_j(k)$, then $x_i(k+1) = x_i(k)$;
(R2) if $x_i(k) < x_j(k)$, then $x_i(k+1) \in (x_i(k), x_j(k)]$;
(R3) if $x_i(k) > x_j(k)$, then $x_i(k+1) \in [x_j(k), x_i(k))$.

Let the subset $\mathcal{C}$ of $\mathbb{Z}^n$ be the set of general consensus states given by $\mathcal{C} := \{x : x_1 = \cdots = x_n\}$. If $x(k) = x^* \in \mathcal{C}$, then $x(k') = x^* \in \mathcal{C}$ for all $k' > k$ because only (R1) applies. We say that the nodes achieve general consensus almost surely if for each $x(0)$, $\Pr[(\exists T < \infty, x^* \in \mathcal{C})(\forall k \geq T)x(k) = x^*] = 1$, where the probability is with respect to the sequence of edges chosen in the algorithm. Under **QC** algorithm, a necessary and sufficient graphical condition that ensures almost sure general consensus is that the digraph $\mathcal{G}$ contains a *globally reachable node* (i.e., a node that is connected to every other node via a directed path) (Cai & Ishii, 2011b). Clearly if $\mathcal{G}$ is complete, then every node is globally reachable.

The *convergence time* of **QC** algorithm is the random variable T_{qc} defined by $T_{qc} := \inf\{k \geq 0 : x(k) \in \mathcal{C}\}$. The mean convergence time (with respect to the probability distribution on edge activation), starting from a state $x_0 \in \mathcal{X}$, is then given by

$$E_{qc}(x_0) := \mathbb{E}\left[T_{qc} | x(0) = x_0\right]. \tag{2}$$

Problem 1. Let Assumptions 1 and 2 hold. Find an upper bound of the mean convergence time $E_{qc}(x_0)$ of **QC** algorithm with respect to all initial states $x_0 \in \mathcal{X}$.

We now present the main result of this section.

Theorem 3. *Let Assumptions 1 and 2 hold. Then* $\max_{x_0\in\mathcal{X}} E_{qc}(x_0) < n(n-1)(M-m) = O(n^2)$.

To derive this bound, we first provide preliminaries on the *hitting time* in finite Markov chains.

2.2. Preliminaries on hitting time

Let $\{X_k\}_{k\geq 0}$ be a Markov chain with a finite state space $\mathcal{S}$ and a transition probability matrix $P = (P_{ij})$ (e.g., Norris (1997)). The entry P_{ij} denotes the one-step transition probability from state i to state j. In particular, the diagonal entry P_{ii} denotes the *selfloop* transition probability. A state $i \in \mathcal{S}$ is said to be *absorbing* if $P_{ii} = 1$. For a given $\{X_k\}_{k\geq 0}$, the *hitting time* of a subset $\mathcal{T}$ of $\mathcal{S}$ is the random variable $H_{\mathcal{T}}\left(\{X_k\}_{k\geq 0}\right)$ defined by $H_{\mathcal{T}}\left(\{X_k\}_{k\geq 0}\right) := \inf\{l \geq 0 : X_l \in \mathcal{T}\}$. The mean time (with respect to the probability distribution specified by P) taken for the chain, starting from a state $i \in \mathcal{S}$, to hit $\mathcal{T}$ is given by

$$\begin{aligned} E_i &:= \mathbb{E}\left[H_{\mathcal{T}}\left(\{X_k\}_{k\geq 0}\right) | X_0 = i\right] \\ &= \sum_{l=0}^{\infty} l \cdot \mathbb{P}\left[H_{\mathcal{T}}\left(\{X_k\}_{k\geq 0}\right) = l | X_0 = i\right], \end{aligned} \tag{3}$$

where $\mathbb{E}[\cdot|\cdot]$ and $\mathbb{P}[\cdot|\cdot]$ denote the conditional expectation and conditional probability operators, respectively.

Using a standard fact on mean hitting time (Norris, 1997, Theorem 1.3.5), we derive a closed-form expression of the mean hitting time for a specific Markov chain; this chain will be shown to characterize the state transition structure under **QC** algorithm. For the proof, refer to Cai and Ishii (2010).