Automatica 70 (2016) 121–127

Contents lists available at ScienceDirect

Automatica

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

Brief paper

Distributed stopping for average consensus in undirected graphs via event-triggered strategies☆

Nicolaos E. Manitara, Christoforos N. Hadjicostis [1]

Department of Electrical and Computer Engineering, University of Cyprus, Nicosia 1678, Cyprus

ARTICLE INFO

Article history:
Received 6 March 2015
Received in revised form
2 November 2015
Accepted 15 March 2016
Available online 15 April 2016

Keywords:
Approximate average consensus
Distributed stopping
Finite-time average consensus
Multi-agent systems

ABSTRACT

We develop and analyze two distributed event-triggered linear iterative algorithms that enable the components of a distributed system, each with some initial value, to reach *approximate average consensus* on their initial values, after executing a *finite* number of iterations. Each proposed algorithm provides a criterion that allows the nodes to determine, in a distributed manner, when to terminate because approximate average consensus has been reached, i.e., all nodes have obtained a value that is within a small distance from the average of their initial values. We focus on a distributed system whose underlying topology is captured by an undirected (symmetric) graph, and develop linear iterative strategies with time-varying weights, chosen based on the subset of edges that separate nodes with significantly different values and are considered active at each iteration. In our simulations, we illustrate the proposed algorithms and compare the number of iterations and transmitted values required by the proposed protocols against a previously proposed stopping protocol for approximate average consensus.

1. Introduction

In distributed systems and networks, it is often necessary for all or some of the nodes (system components) to calculate a function of certain parameters that we refer to as initial values. When all the nodes obtain the average of these initial values, they are said to have reached average consensus. Over the last few decades, a variety of distributed algorithms that allow the components to calculate different functions of their initial values have been proposed by the control, communication, and computer science communities (Cortés, 2008; Olfati-Saber, Alex Fax, & Murray, 2007).

Average consensus and, more generally, consensus have received a lot of attention from the control community due to their usage in various emerging distributed control applications, including wireless smart meters (where all nodes have to determine the average demand or consumption of the network Domínguez-García & Hadjicostis, 2010), and multi-agent systems (where all agents communicate with each other in order to coordinate their direction, speed, and position Ren, Beard, & Atkins, 2005). One popular approach to consensus (to some value, not necessarily the average) is based on a linear iterative strategy, where each node in the network repeatedly updates its value to be a weighted sum of its own previous value and the previous values of its neighbors. In particular, previous work has shown that, if the network topology satisfies certain conditions, the weights for the linear iteration can be chosen such that all of the nodes in the network converge asymptotically to the same value (which, under some additional requirements on the weights, can be the average of their initial values Sundaram & Hadjicostis, 2008; Xiao & Boyd, 2004). Another popular approach to the calculation of this *average* value is based on ratio-consensus (Bénézit, Blondel, Thiran, Tsitsiklis, & Vetterli, 2010; Domínguez-García & Hadjicostis, 2010; Kempe, Dobra, & Gehrke, 2003), which simultaneously runs two linear iterations and allows each node to asymptotically obtain the average as the ratio of the two state variables it maintains.

The above popular approaches use simple local rules to distributively calculate important quantities like the average. The

☆ The material in this paper was partially presented at the 51st Annual Allerton Conference on Communication, Control, and Computing, October 2–4, 2013, Monticello, IL, USA and at the 52nd Annual Allerton Conference on Communication, Control, and Computing, October 1–3, 2014, Monticello, IL, USA (Manitara & Hadjicostis, 2013, 2014a). This paper was recommended for publication in revised form by Associate Editor Dimos V. Dimarogonas under the direction of Editor Christos G. Cassandras. This version of the paper focuses on undirected graphs and provides complete proofs, examples, and simulations/comparisons against previous work, not provided in these conference papers. It also describes and analyzes novel enhancements to the algorithms (i.e., time-varying Metropolis weights) that allow the nodes to reach approximate average consensus faster.

E-mail addresses: manitara.nicolas@ucy.ac.cy (N.E. Manitara), chadjic@ucy.ac.cy (C.N. Hadjicostis).

[1] Tel.: +357 22 892231; fax: +357 22 895079.

http://dx.doi.org/10.1016/j.automatica.2016.03.030
0005-1098/

main problem in the applicability of these techniques, however, is the fact that convergence to the average is asymptotic. Typically, this is handled by a priori determining a number of finite steps that ensures that the nodes have values sufficiently close to the average, but this requires knowledge of the network and the convergence rate of the iteration. The algorithms proposed in this paper provide an alternative that allows the nodes to approximately calculate the average in *finite* time using algorithms that still rely exclusively on local information. Perhaps more importantly, the paper investigates the topic of distributed stopping, i.e., how the nodes can determine when to terminate their transmissions based on locally available information. This question has received limited attention thus far in the control literature, with the notable exception of Yadav and Salapaka (2007), which we discuss in detail later.

In the proposed algorithms, each node makes a decision on how to update and/or transmit its value, based on the difference between its calculated value and the values it receives from its neighbors. Note that, during the execution of the algorithms, a link or a node that becomes inactive at a particular time step may be triggered to become active at a later time step if the node value and/or the value of at least one of its neighbors change in a way that increases their absolute difference. In both algorithms, the iterative process (which relies on linear updates with time-varying weights) ends when all nodes cease transmitting their values, in which case they can be shown to have reached approximate average consensus, i.e., the absolute difference between the final value of each node and the exact average of the initial values is smaller than an error bound, whose value depends on a parameter ε (small real value) and the diameter D of the underlying undirected graph. Unlike existing approaches (e.g., Yadav & Salapaka, 2007), the proposed strategies lead to time-varying strategies and their analysis requires tools from weak ergodicity of inhomogeneous Markov chains. Note that dealing with digraphs (directed graphs) poses a number of additional difficulties, see Manitara and Hadjicostis (2014b).

2. Background

In distributed systems, we can model the network topology as a directed graph (digraph) $G = \{X, E\}$ where $X = \{1, 2, \ldots, n\}$ is the set of components in the system (nodes) and $E \subseteq X \times X - \{(i, i) \mid i \in X\}$ is the set of directed communication links (edges) (West, 2001). In particular, edge $(i, j) \in E$ if node j can send information to node i. The nodes that can transmit (receive) information to (from) node i are said to be the in-neighbors (out-neighbors) of node i and are represented by the set $\mathcal{N}_i^- = \{j \mid (i, j) \in E\}$ ($\mathcal{N}_i^+ = \{j \mid (i, j) \in E\}$); the number of in-neighbors (out-neighbors) of node i is called the in-degree (out-degree) of node i and is denoted by $\mathcal{D}_i^- = |\mathcal{N}_i^-|$ ($\mathcal{D}_i^+ = |\mathcal{N}_i^+|$). In an undirected graph (i.e., a graph for which $(i, j) \in E$ if and only if $(j, i) \in E$), the in-neighbors of node i are identical to its out-neighbors. Thus, we refer to the neighbors $\mathcal{N}_i = \mathcal{N}_i^- = \mathcal{N}_i^+$ and the degree $\mathcal{D}_i = \mathcal{D}_i^- = \mathcal{D}_i^+$ of node $i \in X$.

A path of length t from node j to node i, $i \neq j$, is a sequence of nodes $j = i_0, i_1, \ldots, i_{t-1}, i_t = i$, such that $(i_l, i_{l-1}) \in E$ for all $l = 1, 2, \ldots, t$. The minimum distance from node j to node i, $i \neq j$, is the length of the shortest path from node j to node i; it is denoted by $d_{\min}(i, j)$ and it is taken to be infinite if there is no path from node j to node i. The graph is *strongly connected* (or simply *connected* in the case of an undirected graph) if there exists a path (of finite length) from each node j to each other node i. By convention, $d_{\min}(i, i) = 0$ for all $i \in X$. The diameter D of graph $G = \{X, E\}$ is defined as the longest shortest path between any two nodes, i.e., $D = \max_{i,j \in X, i \neq j} d_{\min}(i, j)$.

2.1. Average consensus via a linear iterative strategy

In average consensus problems, we are given a graph $G = \{X, E\}$ that captures the topology of the distributed system. Each node i has an initial value V_i and the objective is to calculate the average of these initial values, which we denote by $\bar{x} = \frac{1}{n} \sum_{l=1}^{n} V_l$ (where $n = |X|$ is the cardinality of the set X). We assume a broadcast model where each node sends to *all* of its neighbors the same value. Suppose that, at each time-step k, each node updates its value as a weighted sum of its own value and the values of its in-neighbors so that (Jadbabaie, Lin, & Stephen Morse, 2003; Sundaram & Hadjicostis, 2008)

$$x_i[k+1] = p_{ii} x_i[k] + \sum_{j \in \mathcal{N}_i^-} p_{ij} x_j[k], \tag{1}$$

where p_{ij} form a set of (fixed) weights and $x_i[0] = V_i$. The values for all the nodes at time-step k can be aggregated into the value vector $x[k] = [x_1[k], x_2[k], \ldots, x_n[k]]^T$ (where ^{T} denotes matrix/vector transposition) and the update strategy for the entire network can be written compactly as

$$x[k+1] = Px[k], \quad k \in \mathbb{N}, \tag{2}$$

where $x[0] = [V_1, V_2, \ldots, V_n]^T$ and matrix $P = [p_{ij}]$, with the weights p_{ij} satisfying $p_{ij} = 0$ if $j \notin \mathcal{N}_i^- \cup \{i\}$.

Theorem 1 (*Xiao & Boyd, 2004*). *Iteration (2) reaches asymptotic consensus on the linear functional $c^T x[0]$ for some column vector c (under the technical condition that c is normalized so that $c^T \mathbf{1} = 1$ where $\mathbf{1} = [1, 1, \ldots, 1]^T$ is the all ones column vector) if and only if the weight matrix P satisfies the conditions below:*

(1) *P has a simple eigenvalue at 1, with left eigenvector c^T and right eigenvector $\mathbf{1}$;*

(2) *All other eigenvalues of P have magnitude strictly less than 1.*

In particular, if $c = \frac{1}{n}\mathbf{1}$, then average consensus is reached. Also note that if p_{ij} are restricted to be nonnegative, then the above conditions for asymptotic average consensus are equivalent to P being a primitive doubly stochastic matrix.

2.2. Average consensus via ratio consensus

The ratio consensus algorithm performs in parallel two iterative computations (of the type in (1)) and allows each node to asymptotically obtain the exact average of the initial values as the ratio of the two state variables that it maintains. More specifically, each node i maintains, at iteration k, state variables $y_i[k]$ and $z_i[k]$, and updates them, for $k \in \mathbb{N}$, as

$$y_i[k+1] = \sum_{j \in \mathcal{N}_i^- \cup \{i\}} y_j[k] / (1 + \mathcal{D}_j^+), \tag{3}$$

$$z_i[k+1] = \sum_{j \in \mathcal{N}_i^- \cup \{i\}} z_j[k] / (1 + \mathcal{D}_j^+), \tag{4}$$

with $y_i[0] = V_i$ and $z_i[0] = 1$, for $i \in X$. At each time step k, each node i also calculates the ratio $r_i[k] = y_i[k]/z_i[k]$; under the assumption that the digraph describing the exchange of information is strongly connected, it can be shown that $r_i[k]$ asymptotically converges to the average of the initial values (Domínguez-García & Hadjicostis, 2010). Specifically, with the chosen initial conditions, we have that

$$\lim_{k \to \infty} r_i[k] = \frac{\sum_l y_l[0]}{\sum_l z_l[0]} = \frac{\sum_l V_l}{n}, \quad \forall i \in X. \tag{5}$$