Automatica 81 (2017) 221–231

Contents lists available at ScienceDirect

Automatica

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

Differentially private average consensus: Obstructions, trade-offs, and optimal algorithm design☆

Erfan Nozari [a], Pavankumar Tallapragada [b], Jorge Cortés [a]

[a] Department of Mechanical and Aerospace Engineering, University of California, San Diego, United States
[b] Department of Electrical Engineering, Indian Institute of Science, Bengaluru, India

ARTICLE INFO

Article history:
Received 30 December 2015
Received in revised form
3 February 2017
Accepted 28 February 2017

Keywords:
Average consensus
Differential privacy
Multi-agent systems
Exponential mean-square convergence rate
Networked control systems

ABSTRACT

This paper studies the multi-agent average consensus problem under the requirement of differential privacy of the agents' initial states against an adversary that has access to all the messages. We first establish that a differentially private consensus algorithm cannot guarantee convergence of the agents' states to the exact average in distribution, which in turn implies the same impossibility for other stronger notions of convergence. This result motivates our design of a novel differentially private Laplacian consensus algorithm in which agents linearly perturb their state-transition and message-generating functions with exponentially decaying Laplace noise. We prove that our algorithm converges almost surely to an unbiased estimate of the average of agents' initial states, compute the exponential mean-square rate of convergence, and formally characterize its differential privacy properties. We show that the optimal choice of our design parameters (with respect to the variance of the convergence point around the exact average) corresponds to a one-shot perturbation of initial states and compare our design with various counterparts from the literature. Simulations illustrate our results.

1. Introduction

The social adoption of new technologies in networked cyber-physical systems relies heavily on the privacy preservation of individual information. Social networking, the power grid, and smart transportation are a few examples of domains in need of privacy-aware design of control and coordination strategies. In these scenarios, the ability of a networked system to fuse information, compute common estimates of unknown quantities, and agree on a common view of the world is critical. Motivated by these observations, this paper studies the multi-agent average consensus problem, where a group of agents seek to agree on the average of their individual values by only interchanging information with their neighbors. This problem has numerous applications in synchronization, network management, and distributed control/computation/optimization. In the context of privacy preservation, the notion of differential privacy has gained significant popularity due to its rigorous formulation and proven security properties, including resilience to post-processing and side information, and independence from the model of the adversary. Roughly speaking, a strategy is differentially private if the information of an agent has no significant effect on the aggregate output of the algorithm, and hence its data cannot be inferred by an adversary from its execution. This paper is a contribution to the emerging body of research that studies privacy preservation in cooperative network systems, specifically focused on gaining insight into the achievable trade-offs between privacy and performance in multi-agent average consensus.

Literature Review: The problem of multi-agent average consensus has been a subject of extensive research in networked systems and it is impossible to survey here the vast amount of results in the literature. We provide (Bullo, Cortés, & Martínez, 2009; Mesbahi & Egerstedt, 2010; Olfati-Saber, Fax, & Murray, 2007; Ren & Beard, 2008) and the references therein as a starting point for the interested reader. In cyberphysical systems, privacy at the physical layer provides protection beyond the use of higher-level encryption-based techniques. Information-theoretic approaches to privacy at the physical layer have been actively pursued (Gündüz, Erkip, & Poor, 2010; Mukherjee, Fakoorian, Huang, & Swindlehurst, 2014). Recently, these ideas have also been utilized in the context of

☆ The material in this paper was presented at the 5th IFAC Workshop on Distributed Estimation and Control in Networked Systems, September 10–11, 2015, Philadelphia, PA, USA as Nozari, Tallapragada, and Cortés (2014). This paper was recommended for publication in revised form by Associate Editor Claudio De Persis under the direction of Editor Christos G. Cassandras.

E-mail addresses: enozari@ucsd.edu (E. Nozari), pavant@ee.iisc.ernet.in (P. Tallapragada), cortes@ucsd.edu (J. Cortés).

http://dx.doi.org/10.1016/j.automatica.2017.03.016
0005-1098/

control (Tanaka & Sandberg, 2015). The paper Mukherjee et al. (2014) also surveys the more recent game-theoretic approach to the topic. In computer science, the notion of differential privacy, first introduced in Dwork (2006) and Dwork, McSherry, Nissim, and Smith (2006), and the design of differentially private mechanisms have been widely studied in the context of privacy preservation of databases. The work Dwork and Roth (2014) provides a recent comprehensive treatment. A well-known advantage of differential privacy over other notions of privacy is its immunity to post-processing and side information, which makes it particularly well-suited for multi-agent scenarios where agents do not fully trust each other and/or the communication channels are not fully secure. While secure multi-party computation also deals with scenarios where no trust exists among agents, the maximum number of agents that can collude (without the privacy of others being breached) is bounded, whereas using differential privacy provides immunity against arbitrary collusions (Kairouz, Oh, & Viswanath, 2015; Pettai & Laud, 2015). As a result, differential privacy has been adopted by recent works in a number of areas pertaining to networked systems, such as control (Huang, Mitra, & Dullerud, 2012; Huang, Wang, Mitra, & Dullerud, 2014; Wang, Huang, Mitra, & Dullerud, 2014), estimation (Ny & Pappas, 2014), and optimization (Han, Topcu, & Pappas, 2014; Huang, Mitra, & Vaidya, 2015; Nozari, Tallapragada, & Cortés, in press). Of relevance to our present work, the paper Huang et al. (2012) studies the average consensus problem with differential privacy guarantees and proposes an adjacency-based distributed algorithm with decaying Laplace noise and mean-square convergence. The algorithm preserves the differential privacy of the agents' initial states but the expected value of its convergence point depends on the network topology and may not be the exact average, even in expectation. By contrast, the algorithm proposed in this work enjoys almost sure convergence, asymptotic unbiasedness, and an explicit characterization of its convergence rate. Our results also allow individual agents to independently choose their level of privacy. The problem of privacy-preserving average consensus has also been studied using other notions of privacy. The work Manitara and Hadjicostis (2013) builds on Kefayati, Talebi, Khalaj, and Rabiee (2007) to let agents have the option to add to their first set of transmitted messages a zero-sum noise sequence with finite random length, which in turn allows the coordination algorithm to converge to the exact average of their initial states. As long as an adversary cannot listen to the transmitted messages of an agent as well as all its neighbors, the privacy of that agent is preserved, in the sense that different initial conditions may produce the same transmitted messages. This idea is further developed in Mo and Murray (2014, 2017), where agents add an infinitely-long exponentially-decaying zero-sum sequence of Gaussian noise to their transmitted messages. The algorithm has guaranteed mean-square convergence to the average of the agents' initial states and preserves the privacy of the nodes whose messages and those of their neighbors are not listened to by the malicious nodes, in the sense that the maximum-likelihood estimate of their initial states has nonzero variance. Finally, Duan, He, Cheng, Mo, and Chen (2015) considers the problem of privacy preserving maximum consensus.

Statement of Contributions: We study the average consensus problem where a group of agents seek to compute and agree on the average of their local variables while seeking to keep them differentially private against an adversary with potential access to all group communications. This privacy requirement also applies to the case where each agent wants to keep its initial state private against the rest of the group (e.g., due to the possibility of communication leakages). The main contributions of this work are the characterization and optimization of the fundamental trade-offs between differential privacy and average consensus. Our first contribution is the formulation and formal proof of a general impossibility result. We show that as long as a coordination algorithm is differentially private, it is impossible to guarantee the convergence of agents' states to the average of their initial values, even in distribution. This result automatically implies the same impossibility result for stronger notions of convergence. Motivated by it, our second contribution is the design of a linear Laplacian-based consensus algorithm that achieves average consensus in expectation — the most that one can expect. We prove the almost sure convergence and differential privacy of our algorithm and characterize its accuracy and convergence rate. Our final contribution is the computation of the optimal values of the design parameters to achieve the most accurate consensus possible. Letting the agents fix a (local) desired value of the privacy requirement, we minimize the variance of the algorithm convergence point as a function of the noise-to-state gain and the amplitude and decay rate of the noise. We show that the minimum variance is achieved by the one-shot perturbation of the initial states by Laplace noise. This result reveals the optimality of one-shot perturbation for static average consensus, previously (but implicitly) shown in the sense of information-theoretic entropy. Various simulations illustrate our results.

2. Preliminaries

This section introduces notations and basic concepts. We denote the set of reals, positive reals, non-negative reals, positive integers, and nonnegative integers by $\mathbb{R}$, $\mathbb{R}_{>0}$, $\mathbb{R}_{\geq 0}$, $\mathbb{N}$, and $\mathbb{Z}_{\geq 0}$, respectively. We denote by $\|\cdot\|$ the Euclidean norm. We let $(\mathbb{R}^n)^{\mathbb{N}}$ denote the space of vector-valued sequences in $\mathbb{R}^n$. For $\{x(k)\}_{k=0}^{\infty} \in (\mathbb{R}^n)^{\mathbb{N}}$, we use the shorthand notation $\mathbf{x} = \{x(k)\}_{k=0}^{\infty}$ and $\mathbf{x}_k = \{x(j)\}_{j=0}^{k}$. $I_n \in \mathbb{R}^{n\times n}$ and $\mathbf{1}_n \in \mathbb{R}^n$ denote the identity matrix and the vector of ones, respectively. For $x \in \mathbb{R}^n$, $\mathrm{Ave}(x) = \frac{1}{n}\mathbf{1}_n^T x$ denotes the average of its components. We let $\Pi_n = \frac{1}{n}\mathbf{1}_n\mathbf{1}_n^T$. Note that Π_n is diagonalizable, and has one eigenvalue equal to 1 with eigenspace

$$\mathbb{R}\mathbf{1}_n \triangleq \{a\mathbf{1}_n \mid a \in \mathbb{R}\},$$

and all other eigenvalues equal to 0. For a vector space $V \subset \mathbb{R}^n$, we let $V^{\perp}$ denote the vector space orthogonal to V. A matrix $A \in \mathbb{R}^{n\times n}$ is stable if all its eigenvalues have magnitude strictly less than 1. A function $\gamma : [0,\infty) \to [0,\infty)$ belongs to class $\mathcal{K}$ if it is continuous and strictly increasing and $\gamma(0) = 0$. Similarly, a function $\beta : [0,\infty)\times[0,\infty) \to [0,\infty)$ belongs to class $\mathcal{KL}$ if $\beta(\cdot, s)$ belongs to class $\mathcal{K}$ for any $s \in [0,\infty)$ and $\beta(r,\cdot)$ is decreasing and $\lim_{s\to\infty}\beta(r,s) = 0$ for any $r \in [0,\infty)$. For $q \in (0,1)$, the Euler function is given by $\varphi(q) = \prod_{k=1}^{\infty}(1-q^k) > 0$. Note that

$$\lim_{k\to\infty}\prod_{j=k}^{\infty}(1-q^j) = \lim_{k\to\infty}\frac{\varphi(q)}{\prod_{j=1}^{k-1}(1-q^j)} = 1.$$

For a function $f : X \to Y$ and sets $A \subseteq X$ and $B \subseteq Y$, we use $f(A) = \{f(x) \in Y | x \in A\}$ and $f^{-1}(B) = \{x \in X | f(x) \in B\}$. In general, $f(f^{-1}(B)) \subseteq B$. Finally, for any topological space X, we denote by $\mathcal{B}(X)$ the set of Borel subsets of X.

2.1. Graph theory

We present some useful notions on algebraic graph theory following Bullo et al. (2009). Let $\mathcal{G} = (V, E, A)$ denote a weighted undirected graph with vertex set V of cardinality n, edge set $E \subset V \times V$, and symmetric adjacency matrix $A \in \mathbb{R}_{\geq 0}^{n\times n}$. A path from i to j is a sequence of vertices starting from i and ending in j such that any pair of consecutive vertices is an edge. The set of neighbors $\mathcal{N}_i$ of i is the set of nodes j such that $(i,j) \in E$. A graph is connected if for each node there exists a path to any other node. The weighted degree matrix is the diagonal matrix $D \in \mathbb{R}^{n\times n}$ with diagonal $A\mathbf{1}_n$. The Laplacian is $L = D - A$ and has the following properties: