Available online at www.sciencedirect.com

ScienceDirect

IFAC Papers Online
CONFERENCE PAPER ARCHIVE

IFAC-PapersOnLine 48-22 (2015) 203–208

Differentially Private Average Consensus with Optimal Noise Selection

Erfan Nozari [*] Pavankumar Tallapragada [*] Jorge Cortés [*]

[*] *Department of Mechanical and Aerospace Engineering, University of California, San Diego, {enozari,ptallapragada,cortes}@ucsd.edu*

Abstract: This paper studies the problem of privacy-preserving average consensus in multi-agent systems. The network objective is to compute the average of the initial agent states while keeping these values differentially private against an adversary that has access to all inter-agent messages. We establish an impossibility result that shows that exact average consensus cannot be achieved by any algorithm that preserves differential privacy. This result motives our design of a differentially private discrete-time distributed algorithm that corrupts messages with Laplacian noise and is guaranteed to achieve average consensus in expectation. We examine how to optimally select the noise parameters in order to minimize the variance of the network convergence point for a desired level of privacy.

Keywords: average consensus, differential privacy, multi-agent systems

1. INTRODUCTION

Multi-agent average consensus is a basic distributed control 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 found numerous applications in sensor networks, synchronization, network management, and distributed computation and optimization. In many of these applications, guaranteeing the privacy of the individual agents is an important aspect that has not been sufficiently studied in the context of networked systems and cooperative strategies. An increasing number of works look at the notion of differential privacy, which specifies that the information of an agent has no significant effect in the aggregate output of the algorithm, and hence its data cannot be inferred by an adversary from its execution. This is a strong notion of privacy with a rigorous formulation and proven security properties, including resilience to post-processing and auxiliary information and independence from the model of the adversary. This paper is a contribution to this body of research where we focus our attention on gaining insight into the achievable trade-offs between privacy and performance in multi-agent average consensus.

Literature Review There is a large literature on the (average) consensus problem in networked systems and the interested reader is referred to (Bullo et al., 2009; Ren and Beard, 2008; Mesbahi and Egerstedt, 2010) and references therein for a comprehensive review. The notion of differential privacy, first introduced in (Dwork et al., 2006; Dwork, 2006), has been the subject of extensive research in the database literature over the past decade. A recent comprehensive text can be found in (Dwork and Roth, 2014). Recently, this notion has found its way into a number of areas pertaining networked systems including control (Huang et al., 2012, 2014; Wang et al., 2014), estimation (Ny and Pappas, 2014), and optimization (Han et al., 2014; Huang et al., 2015). Of particular relevance to our paper is the work of Huang et al. (2012), which considers the multi-agent average consensus problem and proposes an adjacency-based distributed algorithm with decaying Laplacian noise in the inter-agent messages. The algorithm is differentially private and agents asymptotically agree on a value that may not be the average of their initial states, even in expectation. Our present work improves upon (Huang et al., 2012) by providing a performance bound that sheds light on what can be achieved in terms of differential privacy for general average consensus dynamics and studying a stronger notion of convergence. Our results also allow individual agents to independently choose their level of privacy. Other works have looked at the average consensus problem employing different notions of privacy. Manitara and Hadjicostis (2013) improve upon (Kefayati et al., 2007) to propose a distributed algorithm where any agent has the option to add a zero-sum noise sequence with finite random length to its first set of transmitted messages. Since the sequence is zero-sum, agents converge to the true average. Privacy of a participating agent, understood as the property that different initial conditions produce the same transmitted messages, is preserved if the malicious nodes cannot listen to it and all its neighbors. The work of Mo and Murray (2014) adds infinite-length exponentially-decaying zero-sum noise sequences to inter-agent messages and formally defines privacy as the inability of a malicious node to perfectly recover the initial state of other nodes via maximum-likelihood estimation. The proposed algorithm is mean-square convergent to the true average and preserves the privacy of nodes whose messages and those of their neighbors are not listened to by the malicious nodes.

Statement of Contributions We consider the multi-agent average consensus problem with privacy preservation requirements on the initial agent states. Our main contributions pertain the understanding of the trade-offs between differential privacy and performance, and can be divided into three groups as follows. Our first contribution is a

2405-8963 © 2015, IFAC (International Federation of Automatic Control) Hosting by Elsevier Ltd. All rights reserved.
Peer review under responsibility of International Federation of Automatic Control.
10.1016/j.ifacol.2015.10.331

general result stating that any distributed coordination algorithm cannot simultaneously be differentially private and guarantee weak convergence of agents to the average of their initial states. Our second contribution is the design of a distributed algorithm that guarantees that the agents converge in expectation to the average of their initial states. Our design uses the classical discrete-time Laplacian-based linear stationary dynamics together with additive Laplacian noise processes. We establish the almost sure convergence, unbiasedness, bounded dispersion, and differential privacy of our design in successive results. Our final contribution pertains to the optimal tuning of the design parameters of the algorithm (specifically, the noise-to-state gain of the system and the amplitude and decay rate of the noise) to minimize the variance of the network convergence point for a desired level of privacy. Various simulations illustrate our results. Most of the proofs are omitted for space reasons and will appear elsewhere.

2. PRELIMINARIES

This section introduces notation 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 let $(\mathbb{R}^n)^{\mathbb{N}}$ denote the space of vector-valued sequences in the Euclidean space $\mathbb{R}^n$. Given n numbers $c_1, \dots, c_n \in \mathbb{R}$, $\mathrm{diag}(c_1, \dots, c_n) \in \mathbb{R}^{n\times n}$ denotes a diagonal matrix with $c_1, \dots, c_n$ on its diagonal. For any $\{x(k)\}_{k=0}^{\infty} \in (\mathbb{R}^n)^{\mathbb{N}}$, we define the shorthand notations $\mathbf{x} = \{x(k)\}_{k=0}^{\infty}$ 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, has one eigenvalue equal to 1 associated with eigenspace

$$\mathcal{D}_n = \{x \in \mathbb{R}^n \mid x_i = \mathrm{Ave}(x),\, i \in \{1, \dots, n\}\},$$

while all other eigenvalues equal 0. For a vector space $V \subset \mathbb{R}^n$, we let $V^\perp$ denote the vector space orthogonal to V. We denote the Euclidean norm in $\mathbb{R}^n$ by $\|\cdot\|$. We say a matrix $A \in \mathbb{R}^{n\times n}$ is stable if all its eigenvalues have magnitude strictly less than 1. For $q \in (0,1)$, the Euler function is given by

$$\varphi(q) = \prod_{k=1}^{\infty}(1-q^k) > 0.$$

Note that

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

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 of the graph. The set of neighbors of i, denoted $\mathcal{N}_i$, is the set of nodes j such that $(i,j) \in E$. The graph $\mathcal{G}$ is connected if for each node there exists a path to any other node. The weighted degree matrix of $\mathcal{G}$ is a diagonal matrix $D \in \mathbb{R}^{n\times n}$ whose ith diagonal element, $i \in \{1, \dots, n\}$, is the sum of the ith row of A. The Laplacian of $\mathcal{G}$ is the symmetric matrix

$$L = D - A,$$

and has the following properties:

- L is positive semi-definite;
- $L\mathbf{1}_n = 0$ and $\mathbf{1}_n^T L = 0$, i.e., 0 is an eigenvalue of L corresponding to the eigenspace $\mathcal{D}_n$;
- $\mathcal{G}$ is connected if and only if $\mathrm{rank}(L) = n-1$, so 0 is a simple eigenvalue of L;
- All eigenvalues of L belong to $[0, 2d_{\max}]$, where $d_{\max}$ is the largest element of D.

For convenience, we define $L_{\mathrm{cpt}} = I_n - \Pi_n$.

2.2 Probability Theory

Here we briefly review basic notions on probability following (Papoulis and Pillai, 2002; Durrett, 2010). Consider a probability space $(\Omega, \Sigma, \mathbb{P})$. If $E, F \in \Sigma$ are two events with $E \subseteq F$, then $\mathbb{P}\{E\} \leq \mathbb{P}\{F\}$. For simplicity, we may sometimes denote events of the type $E_p = \{\omega \in \Omega \mid p(\omega)\}$ by $\{p\}$ where p is a logical statement on the elements of Ω. Clearly, for two statements p and q,

$$(p \Rightarrow q) \Rightarrow (\mathbb{P}\{p\} \leq \mathbb{P}\{q\}). \tag{1}$$

A random variable is a function $X : \Omega \to \mathbb{R}$ such that the inverse image of any open set $B \subseteq \mathbb{R}$ belongs to Σ. For any $N \in \mathbb{R}_{>0}$ and any random variable X with finite expected value μ and finite nonzero variance σ^2, Chebyshev's inequality states that

$$\mathbb{P}\{|X - \mu| \geq N\sigma\} \leq \frac{1}{N^2}.$$

Let for a random variable X, $\mathbb{E}[X]$ and F_X denote its expectation and cumulative distribution function, respectively. Then, a sequence of random variables $\{X_k\}_{k\in\mathbb{Z}_{\geq 0}}$ converges to a random variable X

- almost surely (a.s.) or with probability one if
$$\mathbb{P}\{\lim_{k\to\infty} X_k = X\} = 1;$$
- in mean square (m.s.) if $\mathbb{E}[X_k^2], \mathbb{E}[X^2] < \infty$ for all $k \in \mathbb{Z}_{\geq 0}$ and
$$\lim_{k\to\infty}\mathbb{E}[(X_k - X)^2] = 0;$$
- in probability, if for any $\varepsilon > 0$,
$$\lim_{k\to\infty}\mathbb{P}\{|X_k - X| < \varepsilon\} = 1;$$
- in distribution, or weakly, if for any $x \in \mathbb{R}$ at which F_X is continuous,
$$\lim_{k\to\infty} F_{X_k}(x) = F_X(x).$$

Almost sure convergence and convergence in mean square imply convergence in probability, which itself implies convergence in distribution. Moreover, if $\mathbb{P}\{|X_k| \leq \bar{X}\} = 1$ for all $k \in \mathbb{Z}_{\geq 0}$ and some fixed random variable $\bar{X}$ with $\mathbb{E}[\bar{X}^2] < \infty$, then convergence in probability implies mean square convergence, and if X is a constant, then convergence in distribution implies convergence in probability.

A zero-mean random variable X has Laplace distribution with scale $b \in \mathbb{R}_{>0}$, denoted $X \sim \mathrm{Lap}(b)$, if the pdf of X is

$$f_X(x) = \mathcal{L}(x; b) \triangleq \frac{1}{2b}e^{-\frac{|x|}{b}},$$

for any $x \in \mathbb{R}$. It is easy to see that $|X|$ has an exponential distribution with rate $\lambda = \frac{1}{b}$.