
A Privacy Preserving Randomized Gossip Algorithm via Controlled Noise Insertion

Filip Hanzely^aJakub Konečný^bNicolas Loizou^{c,*}Peter Richtárik^{a,c,d}Dmitry Grishchenko^e^a King Abdullah University of Science and Technology, KSA^b Google, USA^c University of Edinburgh, UK^d Moscow Institute of Physics and Technology, Russia^e Université Grenoble Alpes, France

Abstract

In this work² we present a randomized gossip algorithm for solving the average consensus problem while at the same time protecting the information about the initial private values stored at the nodes. We give iteration complexity bounds for the method and perform extensive numerical experiments.

1 Introduction

In this paper we consider the average consensus (AC) problem. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an undirected connected network with node set $\mathcal{V} = \{1, 2, \dots, n\}$ and edges \mathcal{E} such that $|\mathcal{E}| = m$. Each node $i \in \mathcal{V}$ “knows” a private value $c_i \in \mathbb{R}$. The goal of AC is for every node of the network to compute the average of these values, $\bar{c} \stackrel{\text{def}}{=} \frac{1}{n} \sum_i c_i$, in a distributed fashion. That is, the exchange of information can only occur between connected nodes (neighbours).

The literature on distributed protocols for solving the average consensus problem is vast and has long history [18, 19, 1, 8]. In this work we focus on one of the most popular class of methods for solving the average consensus, the randomized gossip algorithms and propose a gossip algorithm for protecting the information of the initial values c_i , in the case when these may be sensitive. In particular, we develop and analyze a privacy preserving variant of the randomized pairwise gossip algorithm (“randomly pick an edge $(i, j) \in \mathcal{E}$ and then replace the values stored at vertices i and j by their average”) first proposed in [2] for solving the average consensus problem. While we shall not formalize the notion of privacy preservation in this work, it will be intuitively clear that our methods indeed make it harder for nodes to infer information about the private values of other nodes, *which might be useful in practice*.

1.1 Related Work on Privacy Preserving Average Consensus

The introduction of notions of privacy within the AC problem is relatively recent in the literature, and the existing works consider two different ideas. In [7], the concept of differential privacy [3] is used to protect the output value \bar{c} computed by all nodes. In this work, an exponentially decaying Laplacian noise is added to the consensus computation. This notion of privacy refers to protection of

*Corresponding author (n.loizou@sms.ed.ac.uk)

²The full-length paper, which includes a number of additional algorithms and results (including proofs of statements and experiments), is available in [6].

the *final average*, and formal guarantees are provided. A different line of work with a more stricter goal is the design of privacy-preserving average consensus protocols that guarantee protection of the *initial values* c_i of the nodes [16, 14, 15]. In this setting each node should be unable to infer a lot about the initial values c_i of any other node. In the existing works, this is mainly achieved with the clever addition of noise through the iterative procedure that guarantees preservation of privacy and at the same time converges to the exact average. We shall however mention, that none of these works address any specific notion of privacy (no clear measure of privacy is presented) and it is still not clear how the formal concept of differential privacy [3] can be applied in this setting.

1.2 Main Contributions

In this work, we present the first randomized gossip algorithm for solving the Average Consensus problem while at the same time protecting the information about the initial values. To the best of our knowledge, this work is the first which combines the *gossip-asynchronous framework* with the privacy concept of protection of the initial values. Note that all the previously mentioned privacy preserving average consensus papers propose protocols which work on the synchronous setting (all nodes update their values simultaneously).

The convergence analysis of proposed gossip protocol (Algorithm 1) is dual in nature. The dual approach is explained in detail in Section 2. It was first proposed for solving linear systems in [5, 12] and then extended to the concept of average consensus problems in [10, 13]. The dual updates immediately correspond to updates of the primal variables, via an affine mapping.

Algorithm 1 is inspired by the works of [14, 15], and protects the initial values by inserting noise in the process. Broadly speaking, in each iteration, each of the sampled nodes first adds a noise to its current value, and an average is computed afterward. Convergence is guaranteed due to the correlation in the noise across iterations. Each node remembers the noise it added last time it was sampled, and in the following iteration, the previously added noise is first subtracted, and a fresh noise of smaller magnitude is added. Empirically, the protection of initial values is provided by first injecting noise into the system, which propagates across the network, but is gradually withdrawn to ensure convergence to the true average.

2 Technical Preliminaries

Primal and Dual Problems Consider solving the (primal) problem of projecting a given vector $c = x^0 \in \mathbb{R}^n$ onto the solution space of a linear system:

$$\min_{x \in \mathbb{R}^n} \{P(x) \stackrel{\text{def}}{=} \frac{1}{2} \|x - x^0\|^2\} \quad \text{subject to} \quad \mathbf{A}x = b, \quad (1)$$

where $\mathbf{A} \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, $x^0 \in \mathbb{R}^n$. We assume the problem is feasible, i.e., that the system $\mathbf{A}x = b$ is consistent. With the above optimization problem we associate the dual problem

$$\max_{y \in \mathbb{R}^m} D(y) \stackrel{\text{def}}{=} (b - \mathbf{A}x^0)^\top y - \frac{1}{2} \|\mathbf{A}^\top y\|^2. \quad (2)$$

The dual is an unconstrained concave (but not necessarily strongly concave) quadratic maximization problem. It can be seen that as soon as the system $\mathbf{A}x = b$ is feasible, the dual problem is bounded. Moreover, all bounded concave quadratics in \mathbb{R}^m can be written in the as $D(y)$ for some matrix \mathbf{A} and vectors b and x^0 (up to an additive constant).

With any dual vector y we associate the primal vector via an affine transformation, $x(y) = x^0 + \mathbf{A}^\top y$. It can be shown that if y^* is dual optimal, then $x^* = x(y^*)$ is primal optimal. Hence, any dual algorithm producing a sequence of dual variables $y^t \rightarrow y^*$ gives rise to a corresponding primal algorithm producing the sequence $x^t \stackrel{\text{def}}{=} x(y^t) \rightarrow x^*$. See [5, 12] for the correspondence between primal and dual methods.

Randomized Gossip Setup: Choosing \mathbf{A} . In the gossip framework we wish (\mathbf{A}, b) to be an *average consensus (AC)* system.

Definition 1. ([10]) Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an undirected graph with $|\mathcal{V}| = n$ and $|\mathcal{E}| = m$. Let \mathbf{A} be a real matrix with n columns. The linear system $\mathbf{A}x = b$ is an “average consensus (AC) system” for graph \mathcal{G} if $\mathbf{A}x = b$ iff $x_i = x_j$ for all $(i, j) \in \mathcal{E}$.

Algorithm 1 Privacy Preserving Gossip Algorithm via Controlled Noise Insertion

Input: vector of private values $c \in \mathbb{R}^n$; initial variances $\sigma_i^2 \in \mathbb{R}_+$ and variance decrease rate ϕ_i such that $0 \leq \phi_i < 1$ for all nodes i .

Initialize: Set $x^0 = c$; $t_1 = t_2 = \dots = t_n = 0$, $v_1^{-1} = v_2^{-1} = \dots = v_n^{-1} = 0$.

for $t = 0, 1, \dots, k-1$ **do**

1. Choose edge $e = (i, j) \in \mathcal{E}$ uniformly at random
2. Generate $v_i^{t_i} \sim N(0, \sigma_i^2)$ and $v_j^{t_j} \sim N(0, \sigma_j^2)$
3. Set $w_i^{t_i} = \phi_i^{t_i} v_i^{t_i} - \phi_i^{t_i-1} v_i^{t_i-1}$ and $w_j^{t_j} = \phi_j^{t_j} v_j^{t_j} - \phi_j^{t_j-1} v_j^{t_j-1}$
4. Update the primal variable: $x_i^{t+1} = x_j^{t+1} = \frac{x_i^t + w_i^{t_i} + x_j^t + w_j^{t_j}}{2}, \forall l \neq i, j : x_l^{t+1} = x_l^t$
5. Set $t_i = t_i + 1$ and $t_j = t_j + 1$

end

return x^k

In the rest of this paper we focus on a specific AC system; one in which the matrix \mathbf{A} is the incidence matrix of the graph \mathcal{G} (see Model 1 in [5]). In particular, we let $\mathbf{A} \in \mathbb{R}^{m \times n}$ be the matrix defined as follows. Row $e = (i, j) \in \mathcal{E}$ of \mathbf{A} is given by $\mathbf{A}_{ei} = 1$, $\mathbf{A}_{ej} = -1$ and $\mathbf{A}_{el} = 0$ if $l \notin \{i, j\}$. Notice that the system $\mathbf{A}x = 0$ encodes the constraints $x_i = x_j$ for all $(i, j) \in \mathcal{E}$, as desired. It is also known that randomized Kaczmarz method [17, 4, 11] applied to Problem 1 is equivalent to randomized gossip algorithm (see [10, 13, 9] for more details).

3 Private Gossip via Controlled Noise Insertion

In this section, we present the Gossip algorithm with Controlled Noise Insertion. As mentioned in the introduction, the approach is similar to the technique proposed in [14, 15]. Those works, however, address only algorithms in the synchronous setting, while our work is the first to use this idea in the asynchronous setting. Unlike the above, we provide finite time convergence guarantees and allow each node to add the noise differently, which yields a stronger result.

In our approach, each node adds noise to the computation independently of all other nodes. However, the noise added is correlated between iterations for each node. We assume that every node owns two parameters — the initial magnitude of the generated noise σ_i^2 and rate of decay of the noise ϕ_i . The node inserts noise $w_i^{t_i}$ to the system every time that an edge corresponding to the node was chosen, where variable t_i carries an information how many times the noise was added to the system in the past by node i . Thus, if we denote by t the current number of iterations, we have $\sum_{i=1}^n t_i = 2t$.

In order to ensure convergence to the optimal solution, we need to choose a specific structure of the noise in order to guarantee the mean of the values x_i converges to the initial mean. In particular, in each iteration a node i is selected, we subtract the noise that was added last time, and add a fresh noise with smaller magnitude: $w_i^{t_i} = \phi_i^{t_i} v_i^{t_i} - \phi_i^{t_i-1} v_i^{t_i-1}$, where $0 \leq \phi_i < 1$, $v_i^{-1} = 0$ and $v_i^{t_i} \sim N(0, \sigma_i^2)$ for all iteration counters $k_i \geq 0$ is independent to all other randomness in the algorithm. This ensures that all noise added initially is gradually withdrawn from the whole network.

After the addition of noise, a standard Gossip update is made, which sets the values of sampled nodes to their average. Hence, we have $\lim_{t \rightarrow \infty} \mathbb{E} \left[\left(\bar{c} - \frac{1}{n} \sum_{i=1}^n x_i^t \right)^2 \right] = 0$, as desired.

It is not the purpose of this paper to define any quantifiable notion of protection of the initial values formally. However, we note that it is likely the case that the protection of private value c_i will be stronger for bigger σ_i and for ϕ_i closer to 1.

We now provide results of dual analysis of Algorithm 1.

Theorem 2. Let us define $\rho \stackrel{\text{def}}{=} 1 - \frac{\alpha(\mathcal{G})}{2m}$ and $\psi^t \stackrel{\text{def}}{=} \frac{1}{\sum_{i=1}^n (d_i \sigma_i^2)} \sum_{i=1}^n d_i \sigma_i^2 \left(1 - \frac{d_i}{m} (1 - \phi_i^2) \right)^t$, where $\alpha(\mathcal{G})$ stands for algebraic connectivity of \mathcal{G} and d_i denotes the degree of node i . Then for all

$k \geq 1$ we have the following bound

$$\mathbb{E} [D(y^*) - D(y^k)] \leq \rho^k (D(y^*) - D(y^0)) + \frac{\sum (d_i \sigma_i^2)}{4m} \sum_{t=1}^k \rho^{k-t} \psi^t.$$

Note that ψ^t is a weighted sum of t -th powers of real numbers smaller than one. For large enough t , this quantity will depend on the largest of these numbers. This brings us to define M as the set of indices i for which the quantity $1 - \frac{d_i}{m} (1 - \phi_i^2)$ is maximized: $M = \arg \max_i \{1 - \frac{d_i}{m} (1 - \phi_i^2)\}$. Then for any $i_{\max} \in M$ we have

$$\psi^t \approx \frac{1}{\sum_{i=1}^n (d_i \sigma_i^2)} \sum_{i \in M} d_i \sigma_i^2 \left(1 - \frac{d_i}{m} (1 - \phi_i^2)\right)^t = \frac{\sum_{i \in M} d_i \sigma_i^2}{\sum_{i=1}^n (d_i \sigma_i^2)} \left(1 - \frac{d_{i_{\max}}}{m} (1 - \phi_{i_{\max}}^2)\right)^t,$$

which means that increasing ϕ_j for $j \notin M$ will not substantially influence convergence rate. Note that as soon as we have

$$\rho > 1 - \frac{d_i}{m} (1 - \phi_i^2) \quad (3)$$

for all i , the rate from theorem 2 will be driven by ρ^k (as $k \rightarrow \infty$) and we will have $\mathbb{E} [D(y^*) - D(y^k)] = \tilde{O}(\rho^k)$. One can think of the above as a threshold: if there is i such that ϕ_i is large enough so that the inequality (3) does not hold, the convergence rate is driven by $\phi_{i_{\max}}$. Otherwise, the rate is not influenced by the insertion of noise. Thus, in theory, we do not pay anything in terms of performance as long as we do not hit the threshold. One might be interested in choosing ϕ_i so that the threshold is attained for all i , and thus $M = \{1, \dots, n\}$. This motivates the following result:

Corollary 3. *Let us choose $\phi_i \stackrel{\text{def}}{=} \sqrt{1 - \frac{\gamma}{d_i}}$ for all i , where $\gamma \leq d_{\min}$. Then*

$$\mathbb{E} [D(y^*) - D(y^k)] \leq \left(1 - \min\left(\frac{\alpha(\mathcal{G})}{2m}, \frac{\gamma}{m}\right)\right)^k \left(D(y^*) - D(y^0) + \frac{\sum_{i=1}^n (d_i \sigma_i^2)}{4m} k\right).$$

As a consequence, $\phi_i = \sqrt{1 - \frac{\alpha(\mathcal{G})}{2d_i}}$ is the largest decrease rate of noise for node i such that the guaranteed convergence rate of the algorithm is not violated.

4 Experiments

In this section we present a preliminary experiment (for more experiments see Section ??, in the Appendix) to evaluate the performance of the Algorithm 1 for solving the Average Consensus problem. The algorithm has two different parameters for each node i . These are the initial variance $\sigma_i^2 \geq 0$ and the rate of decay, ϕ_i , of the noise.

In this experiment we use two popular graph topologies the cycle graph (ring network) with $n = 10$ nodes and the random geometric graph with $n = 100$ nodes and radius $r = \sqrt{\log(n)/n}$.

In particular, we run Algorithm 1 with $\sigma_i = 1$ for all i , and set $\phi_i = \phi$ for all i and some ϕ . We study the effect of varying the value of ϕ on the convergence of the algorithm.

In Figure 1 we see that for small values of ϕ , we eventually recover the same rate of linear convergence as the Standard Pairwise Gossip algorithm (Baseline) of [2]. If the value of ϕ is sufficiently close to 1 however, the rate is driven by the noise and not by the convergence of the Standard Gossip algorithm. This value is $\phi = 0.98$ for cycle graph, and $\phi = 0.995$ for the random geometric graph in the plots we present.

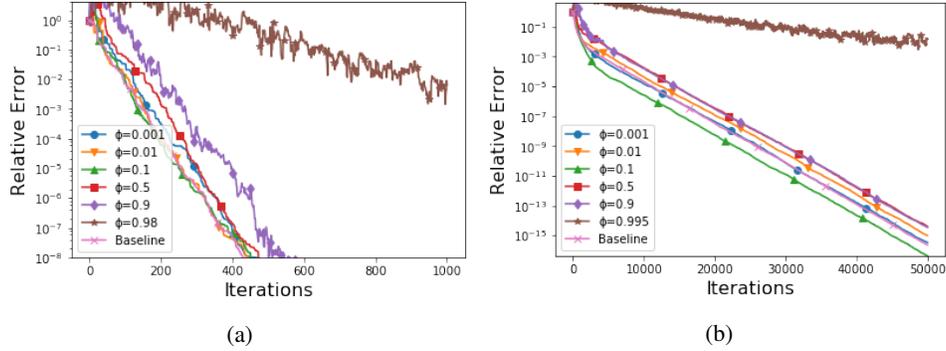


Figure 1: Convergence of Algorithm 1, on the cycle graph (left) and random geometric graph (right) for different values of ϕ . The “Relative Error ” on the vertical axis represents the $\frac{\|x^t - x^*\|^2}{\|x^0 - x^*\|^2}$

References

- [1] Dimitri P Bertsekas and John N Tsitsiklis. *Parallel and distributed computation: numerical methods*, volume 23. Prentice hall Englewood Cliffs, NJ, 1989.
- [2] Stephen Boyd, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 14(SI):2508–2530, 2006.
- [3] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- [4] Robert M Gower and Peter Richtárik. Randomized iterative methods for linear systems. *SIAM Journal on Matrix Analysis and Applications*, 36(4):1660–1690, 2015.
- [5] Robert M Gower and Peter Richtárik. Stochastic dual ascent for solving linear systems. *arXiv:1512.06890*, 2015.
- [6] F. Hanzely, J. Konečný, N. Loizou, P. Richtárik, and D. Grishchenko. Privacy preserving randomized gossip algorithms. *arXiv:1706.07636*, 2017.
- [7] Zhenqi Huang, Sayan Mitra, and Geir Dullerud. Differentially private iterative synchronous consensus. In *Proceedings of the 2012 ACM Workshop on Privacy in the Electronic Society*, pages 81–90. ACM, 2012.
- [8] David Kempe, Alin Dobra, and Johannes Gehrke. Gossip-based computation of aggregate information. In *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on*, pages 482–491. IEEE, 2003.
- [9] Nicolas Loizou, Michael Rabbat, and Peter Richtárik. Provably accelerated randomized gossip algorithms. *arXiv:1810.13084*, 2018.
- [10] Nicolas Loizou and Peter Richtárik. A new perspective on randomized gossip algorithms. In *4th IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, 2016.
- [11] Nicolas Loizou and Peter Richtárik. Linearly convergent stochastic heavy ball method for minimizing generalization error. *NIPS Optimization for Machine Learning Workshop*, 2017.
- [12] Nicolas Loizou and Peter Richtárik. Momentum and stochastic momentum for stochastic gradient, newton, proximal point and subspace descent methods. *arXiv:1712.09677*, 2017.
- [13] Nicolas Loizou and Peter Richtárik. Accelerated gossip via stochastic heavy ball method. *arXiv:1809.08657*, 2018.
- [14] Nicolaos E Maniara and Christoforos N Hadjicostis. Privacy-preserving asymptotic average consensus. In *Control Conference (ECC), 2013 European*, pages 760–765. IEEE, 2013.
- [15] Yilin Mo and Richard M Murray. Privacy preserving average consensus. *IEEE Transactions on Automatic Control*, 62(2):753–765, 2017.
- [16] Erfan Nozari, Pavankumar Tallapragada, and Jorge Cortés. Differentially private average consensus: obstructions, trade-offs, and optimal algorithm design. *Automatica*, 81:221–231, 2017.

- [17] Thomas Strohmer and Roman Vershynin. A randomized Kaczmarz algorithm with exponential convergence. *Journal of Fourier Analysis and Applications*, 15(2):262–278, 2009.
- [18] John N Tsitsiklis. Problems in decentralized decision making and computation. Technical report, DTIC Document, 1984.
- [19] John N Tsitsiklis, Dimitri Bertsekas, and Michael Athans. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE Transactions on Automatic Control*, 31(9):803–812, 1986.