A Privacy Preserving Randomized Gossip Algorithm via Controlled Noise Insertion

Filip Hanzely[1] Jakub Konečný[2] Nicolas Loizou[3] Peter Richtárik[1, 3, 4] Dmitry Grishchenko[5]

[1]KAUST [2]Google [3]University of Edinburgh [4]MIPT [5]Université Grenoble Alpes

Average Consensus Problem (ACP)

SETUP: $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a connected graph with $|\mathcal{V}| = n$ nodes (e.g., sensors) and $|\mathcal{E}| = m$ edges (e.g., communication links). Node $i \in \mathcal{V}$ stores a private value $c_i \in \mathbb{R}$ (e.g., temperature). Let $\alpha(\mathcal{G})$ be the algebraic connectivity of $\mathcal{G}$.

GOAL: Compute the average of the private values (i.e., the quantity $\bar{c} := \frac{1}{n}\sum_i c_i$) in a **distributed** fashion. That is, exchange of information can only occur along the edges.

Primal and Dual Problems

Observe that the optimal solution of the problem

$$\min_{x \in \mathbb{R}^n} \quad \frac{1}{2}\sum_i (x_i - c_i)^2 \quad \text{s.t} \quad x_i = x_j \quad \forall \quad (i,j) \in \mathcal{E} \tag{1}$$

is $x_i^* = \bar{c}$ for all i. The constraints can be written as $\mathbf{A}x = 0$, where $\mathbf{A} \in \mathbb{R}^{m \times n}$, and the rows of $\mathbf{A}$ enforce the constraints $x_i = x_j$ for $(i,j) \in \mathcal{E}$.

Primal Problem (equivalent to (1))

Consider solving the (primal) problem of projecting a given vector $c \in \mathbb{R}^n$ of (private values) onto the solution space of a linear system:

$$\min_{x \in \mathbb{R}^n} \frac{1}{2}\|x - c\|^2 \quad \text{subject to} \quad \mathbf{A}x = 0, \tag{2}$$

where $\mathbf{A} \in \mathbb{R}^{m \times n}$.

Dual Problem

$$\max_{y \in \mathbb{R}^m} D(y) \stackrel{\text{def}}{=} -c^\top \mathbf{A}^\top y - \frac{1}{2}\|\mathbf{A}^\top y\|^2. \tag{3}$$

This is an unconstrained concave quadratic maximization problem.

Lemma: Primal-Dual Relationship

Suppose that $x \in \mathbb{R}^n$ is the primal variable corresponding to the dual variable $y \in \mathbb{R}^m$ through the affine mapping $x \leftarrow c + \mathbf{A}^\top y$. Then:

$$D(y^*) - D(y) = \frac{1}{2}\|x - x^*\|^2,$$

where y^* is any solution of the dual problem.

Randomized Gossip

Algorithm 1 Pairwise Randomized Gossip [1, 2]

Parameters: 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 $0 \le \phi_i < 1$ for all nodes i.
Initialize: Set $x^0 = c$
for $t = 0, 1, \ldots k-1$ **do**
  Choose edge $e = (i,j) \in \mathcal{E}$ uniformly at random
  Update the primal variable: $x_l^{t+1} = \begin{cases} \frac{x_i^t + x_j^t}{2}, & l \in \{i,j\} \\ x_l^t, & l \notin \{i,j\}. \end{cases}$
end for
return x^k

Theorem [1, 2]

The random iterates of the randomized gossip algorithm converge to $x^* = (\bar{c}, \cdots, \bar{c})^\top$, where $\bar{c} = \frac{1}{n}\sum_i c_i$, at a linear rate:

$$\mathbf{E}\left[\|x^k - x^*\|^2\right] \le \left(1 - \frac{\alpha(\mathcal{G})}{2m}\right)^k \|x^0 - x^*\|^2,$$

where $\alpha(\mathcal{G})$ is the algebraic connectivity of graph $\mathcal{G}$.

Privacy Preserving Randomized Gossip

Algorithm 2 Privacy Preserving Randomized Gossip via Controlled Noise Insertion

Parameters: 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 $0 \le \phi_i < 1$ for all nodes i.
Initialize: Set $x^0 = c$; $t_1 = t_2 = \cdots = t_n = 0$, $v_1^{-1} = v_2^{-1} = \cdots = v_n^{-1} = 0$.
for $t = 0, 1, \ldots k-1$ **do**
  Choose edge $e = (i,j) \in \mathcal{E}$ uniformly at random
  Generate $v_i^{t_i} \sim N(0, \sigma_i^2)$ and $v_j^{t_j} \sim N(0, \sigma_j^2)$
  Set $w_i^{t_i} = \phi_i^{t_i} v_i^{t_i} - \phi_i^{t_i - 1} v_i^{t_i - 1}$ and $w_j^{t_j} = \phi_j^{t_j} v_j^{t_j} - \phi_j^{t_j - 1} v_j^{t_j - 1}$
  Update the primal variable: $x_i^{t+1} = x_j^{t+1} = \frac{x_i^t + w_i^{t_i} + x_j^t + w_j^{t_j}}{2}$, $\forall\, l \ne i,j:\ x_l^{t+1} = x_l^t$
  Set $t_i = t_i + 1$ and $t_j = t_j + 1$
end for
return x^k

Main Theorem (Linear Convergence in the Dual)

Let $\rho \stackrel{\text{def}}{=} 1 - \frac{\alpha(\mathcal{G})}{2m}$ 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 $k \ge 1$,

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

If we further choose $\phi_i \stackrel{\text{def}}{=} \sqrt{1 - \frac{\gamma}{d_i}}$ for all i, where $\gamma \le d_{\min}$, then

$$\mathbf{E}\left[D(y^*) - D(y^k)\right] \le \left(1 - \min\left(\frac{\alpha(\mathcal{G})}{2m}, \frac{\gamma}{m}\right)\right)^k \left(D(y^*) - D(y^0) + \mathcal{O}(k)\right).$$

Experiments

Fixed variance $\sigma_i = 1$, $\forall i$, identical decay rates $\phi_i = \phi$:

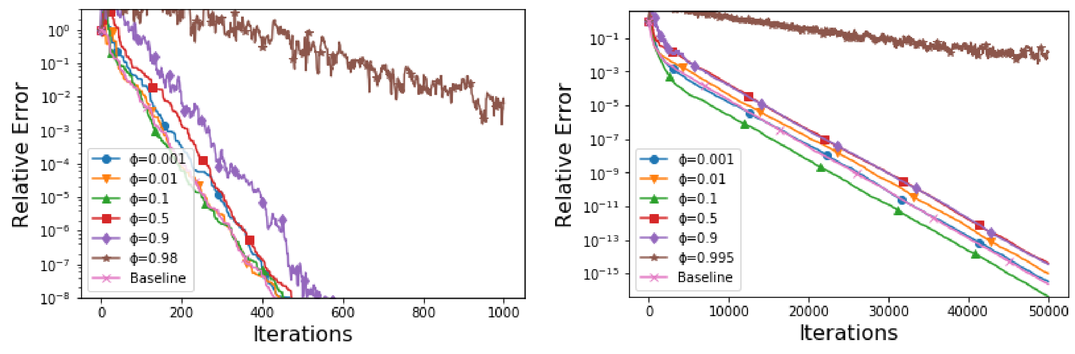

(a) (b)

Figure: Convergence of Algorithm 2, on the cycle graph with 10 nodes (left) and random geometric graph with 100 nodes (right) for different values of ϕ. The "Relative Error " on the vertical axis represents the $\frac{\|x^k - x^*\|^2}{\|x^0 - x^*\|^2}$

Fixed variance $\sigma_i = 1$, $\forall i$, different decay rates $\phi_i = \sqrt{1 - \frac{\gamma}{d_i}}$:

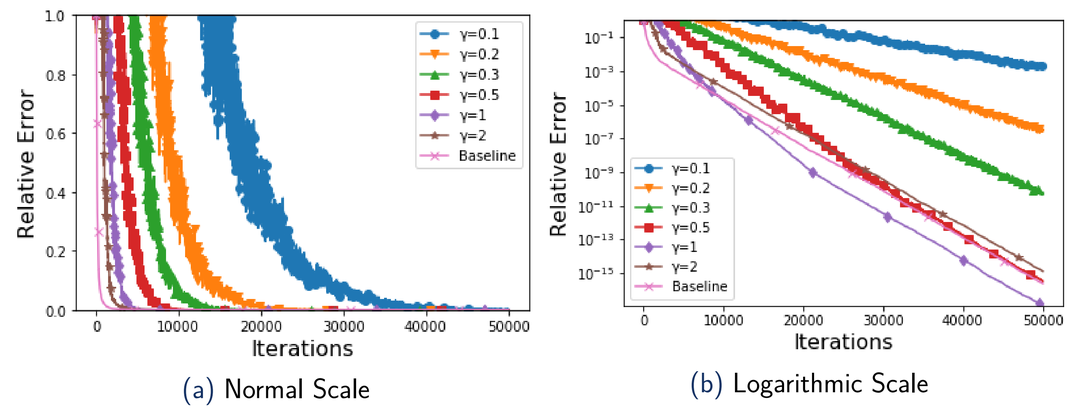

(a) Normal Scale (b) Logarithmic Scale

Figure: Convergence of Algorithm 2 on random geometric graph with 100 nodes for different values of ϕ_i, controlled by γ.

Impact of varying ϕ_i:

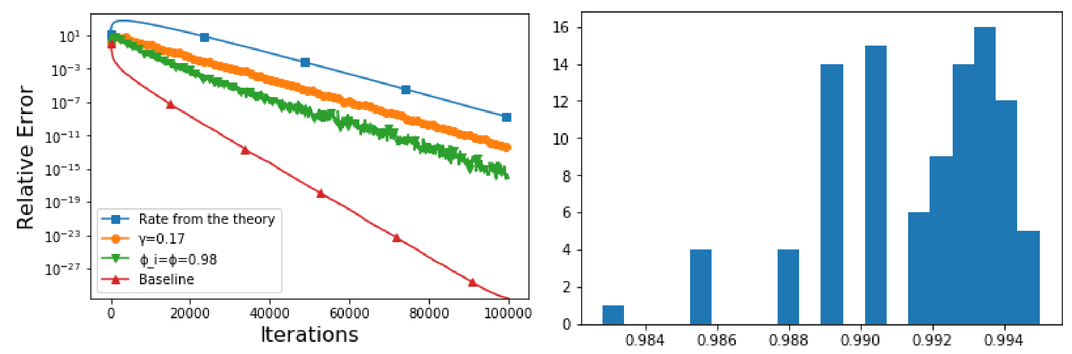

Figure: Network: random geometric graph with 100 nodes. Left: Performance of Algorithm 2 with noise decrease rate chosen according to $\phi_i = \sqrt{1 - \frac{\alpha(\mathcal{G})}{2d_i}}$. Right: Histogram of of distribution of ϕ_i

References

[1] Stephen Boyd, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 14(SI):2508–2530, 2006.

[2] Nicolas Loizou and Peter Richtárik. A new perspective on randomized gossip algorithms. In *4th IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, 2016.

[3] F. Hanzely, J. Konečný, N. Loizou, P. Richtárik, and D. Grishchenko. Privacy preserving randomized gossip algorithms. *arXiv:1706.07636*, 2017.