ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) | 979-8-3503-6874-1/25/$31.00 ©2025 IEEE | DOI: 10.1109/ICASSP49660.2025.10888860

Privacy-Preserving Distributed Maximum Consensus Without Accuracy Loss

Wenrui Yu
CISPA Helmholtz Center for Information Security
Germany
wenrui.yu@cispa.de

Richard Heusdens
Netherlands Defence Academy
Delft University of Technology
the Netherlands
r.heusdens@tudelft.nl

Jun Pang
University of Luxembourg
Luxembourg
jun.pang@uni.lu

Qiongxiu Li
Aalborg University
Denmark
qili@es.aau.dk

***Abstract*—In distributed networks, calculating the maximum element is a fundamental task in data analysis, known as the distributed maximum consensus problem. However, the sensitive nature of the data involved makes privacy protection essential. Despite its importance, privacy in distributed maximum consensus has received limited attention in the literature. Traditional privacy-preserving methods typically add noise to updates, degrading the accuracy of the final result. To overcome these limitations, we propose a novel distributed optimization-based approach that preserves privacy without sacrificing accuracy. Our method introduces virtual nodes to form an augmented graph and leverages a carefully designed initialization process to ensure the privacy of honest participants, even when all their neighboring nodes are dishonest. Through a comprehensive information-theoretical analysis, we derive a sufficient condition to protect private data against both passive and eavesdropping adversaries. Extensive experiments validate the effectiveness of our approach, demonstrating that it not only preserves perfect privacy but also maintains accuracy, outperforming existing noise-based methods that typically suffer from accuracy loss.**

***Index Terms*—maximum consensus, distributed optimization, privacy, information-theoretical analysis, adversary**

I. Introduction

Consensus algorithms are designed to facilitate network-wide agreement through localized computations and the exchange of information among neighboring nodes. These algorithms represent a fundamental challenge in distributed optimization and have found widespread applications. Typical examples include averaging [1], [2], maximum/minimum [3], and median [3] consensus. However, since information sharing is an essential process in solving consensus problems, it raises severe privacy concerns.

Common privacy preservation techniques in consensus problems include differential privacy (DP) [4]–[8], secure multi-party computation (SMPC) [9]–[16], subspace perturbation [17]–[19] and variants of it [20]–[22]. DP achieves a level of protection by adding zero-mean noise, thereby obfuscating the private data. However, this approach involves a tradeoff between utility and privacy; higher levels of noise lead to better privacy but result in reduced accuracy. SMPC techniques, such as secret sharing [23], often incur communication overhead due to the need to split and distribute the message for transmission. Subspace perturbation, based on distributed optimizers such as the Alternating Direction Method of Multipliers (ADMM) [24] or the Primal-Dual Method of Multipliers (PDMM) [25], [26], operates by introducing noise due to proper initialization of the optimization variables. Since algorithms are guaranteed to converge regardless of the initial conditions, the algorithm accuracy remains uncompromised. Consequently, it allows for privacy preservation while maintaining the integrity of the original data.

While average consensus has been extensively studied, the issue of privacy leakage in nonlinear consensus problems, such as maximum/minimum and median consensus, has received relatively little attention. The investigation can advance the understanding of Byzantine robustness in distributed systems, such as federated learning [27]. A few works have attempted to address this concern. Wang et al. [28] directly adds Gaussian noise to private data before broadcasting it to the network, while Venkategowda et al. [29], [30] employs DP within the ADMM framework by adding diminishing noise to the primal variable. Unfortunately, the tradeoff between accuracy and privacy still remains. To overcome it, subspace perturbation [17]–[19] has emerged as an attractive alternative, bypassing it through the proper initialization of auxiliary variables. However, two challenges arise when applying it to maximum consensus. Firstly, this technique was originally proposed for problems with equality constraints, it is unclear whether it works effectively for inequality constraints. Secondly, it guarantees the privacy of an honest node only if it has at least one honest neighbor, which may not always be practical.

In this paper, we propose a simple yet effective approach to address these challenges. Our method not only extends subspace perturbation to inequality constrained scenarios within maximum consensus but also introduces additional virtual nodes (referred to as dummy nodes) to form an augmented graph to ensure the privacy of honest nodes, even in the extreme case that all their neighboring nodes are dishonest. Our approach is grounded by information-theoretical analysis, from which we derive a sufficient condition to ensure (asymptotically) perfect privacy of honest nodes. To our knowledge, it is the first instance of a privacy-preserving maximum consensus algorithm that incurs no accuracy loss while being supported by rigorous information-theoretical analysis. Extensive exper-

This research is partially supported by the EU ChipsJU and the Innovation Fund Denmark through the project CLEVER (no. 101097560).

imental results consolidate the effectiveness of our approach.

II. Preliminaries

A. Problem formulation

We model our network by a graphical model $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{1, \ldots, n\}$ represents the set of nodes/participants in the network and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ represents the set of undirected edges indicating the connections between the nodes (communication links). For each node i we denote its set of neighbors as $\mathcal{N}_i = \{j \in \mathcal{V} \,|\, (i,j) \in \mathcal{E}\}$ and its degree by $d_i = |\mathcal{N}_i|$. Let $s_i \in \mathbb{R}$ denote the data[1] in node $i \in \mathcal{V}$. The privacy-preserving maximum consensus problem is to find the maximum value $s_{\max} = \max\{s_i : i \in \mathcal{V}\}$ in the network without revealing the local data s_i. To do so, we formulate the optimization challenge as a linear programming (LP) problem [29]

$$\begin{aligned} &\text{minimize} && \sum_{i \in \mathcal{V}} x_i, \\ &\text{subject to} && x_i - x_j = 0, \quad (i,j) \in \mathcal{E}, \\ &&& x_i \geq s_i, \quad i \in \mathcal{V}. \end{aligned} \tag{1}$$

When x is updated iteratively, we write $x^{(t)}$ to indicate the update of x at the tth iteration. When we consider x as a realization of a random variable, the corresponding random variable will be denoted by X (corresponding capital).

B. A/PDMM with linear equality and inequality constraints

Following [26], we consider the minimization of a separable convex function subject to a set of inequality constraints by

$$\begin{aligned} &\text{minimize} && \sum_{i \in \mathcal{V}} f_i\left(x_i\right), \\ &\text{subject to} && A_{ij} x_i + A_{ji} x_j \preceq b_{ij}, \quad (i,j) \in \mathcal{E}, \end{aligned} \tag{2}$$

where f_i are convex, closed and proper (CCP) functions and $\preceq$ (generalized inequality) represents element-wise inequality. Constraints between entries are defined by A_{ij}, A_{ji} and b_{ij}.

To solve (2), the update equations of the so-called inequality constraint primal-dual method of multipliers (IEQ-PDMM) [26] for node $i \in \mathcal{V}$ are given by

$$\begin{aligned} x_i^{(t+1)} &= \arg\min_{x_i} \big(f_i(x) \\ &\quad + \sum_{j \in \mathcal{N}_i} \big(z_{i|j}^{(t)} A_{ij} x_i + \frac{c}{2} ||A_{ij} x_i - \frac{1}{2} b_{ij}||^2\big)\big), \\ y_{i|j}^{(t+1)} &= z_{i|j}^{(t)} + 2c(A_{ij} x_i^{(t+1)} - \frac{1}{2} b_{ij}), \\ z_{i|j}^{(t+1)} &= \begin{cases} (1-\theta) z_{i|j}^{(t)} + \theta y_{j|i}^{(t+1)}, & y_{i|j}^{(t+1)} + y_{j|i}^{(t+1)} > 0, \\ (1-\theta) z_{i|j}^{(t)} - \theta y_{i|j}^{(t+1)}, & \text{otherwise,} \end{cases} \end{aligned} \tag{3}$$

where y and z are auxiliary variables, $\theta \in (0,1)$ is an avaraging constant and $c > 0$ is a constant controling the convergence rate. When the objective function is uniformly convex, the algorithm will also converge for $\theta = 1$ (standard PDMM) [26]. Since the LP problem is not uniformly convex, we primarily focus on analyzing the case where $\theta = \frac{1}{2}$. The choice corresponds to the $\frac{1}{2}$-averaged version of PDMM, which is equivalent to ADMM.

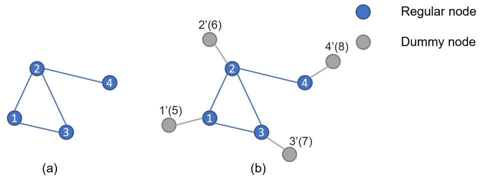

Fig. 1. (a) Example of the original graph $\mathcal{G}$; (b) Example of the augmented graph $\mathcal{G}'$ including dummy nodes.

C. Adversary model and evaluation metrics

Adversary model: We consider two widely used adversary models. The first is the passive adversary, represented by corrupt nodes in the graph. These nodes follow the algorithm's instructions but collude to gather and share information. We denote the set of corrupt nodes in the network by $\mathcal{V}_c$ and the set of honest nodes by $\mathcal{V}_h$. The second type is eavesdropping, which can intercept all messages transmitted through unencrypted channels. These two adversaries are assumed to be able to collaborate to infer the private data of honest nodes.

The performance of the algorithm is evaluated based on the following two requirements and their corresponding metrics.

Output accuracy: It measures how close the optimization results of the privacy-preserving algorithm are to those original non-privacy-preserving algorithms. We quantify the accuracy using the squared error $\|x_i^{(t_{\max})} - x^*\|_2^2$, where $t_{\max}$ denotes the maximum number of iterations and x^* the optimal solution.

Individual privacy: Both ϵ-DP and mutual information approaches are widely used information-theoretical methods for quantifying privacy [31], [32]. We adopt mutual information as the metric for assessing individual privacy as it is shown effective in the literature [33]–[35]. Given the random variable S_i representing the private data at node i and $\mathcal{O}$ representing the total information that the adversary can observe, the mutual information $I(S_i;\mathcal{O})$ [36] measures the amount of information learned about S_i by observing $\mathcal{O}$, which is give by

$$I(S_i;\mathcal{O}) = h(S_i) - h(S_i \mid \mathcal{O}),$$

where $h(\cdot)$ denotes differential entropy. When $I(S_i;\mathcal{O}) = h(S_i)$, the adversary has sufficient information to fully infer s_i. When $I(S_i;\mathcal{O}) = 0$, the adversary cannot infer any information about S_i given the available information $\mathcal{O}$.

III. Proposed approach

We now proceed to the proposed approach. We first introduce how to reformulate the problem using an augmented graph by adding dummy nodes to the network. One for each node, which serves the purpose of overcoming the limitation of requiring at least one honest neighbor for privacy preservation. That is, given node $i \in \mathcal{V}$, we introduce a dummy node i'. The new graph thus obtained will be denoted by $\mathcal{G}' = (\mathcal{V}', \mathcal{E}')$ where $|\mathcal{V}'| = 2|\mathcal{V}|$. See Fig. 1 for an illustration. With this, we can formulate the constraints in (1) as

$$A_{ij} x_i + A_{ji} x_j = 0 \quad \text{for} \quad (i,j) \in \mathcal{E}$$
$$A_{ii'} x_i + A_{i'i} x_j \leq -s_i \quad \text{for} \quad (i,i') \in \mathcal{E}' \backslash \mathcal{E}$$

where $A_{ij} = -A_{ji} = 1$ when $i < j$, and $A_{ii'} = -1$, $A_{i'i} = 0$.

In addition to adding dummy nodes to the graph, we utilize the concept of subspace perturbation, initially introduced for

[1] For simplicity, we assume s_i is a scalar, but results can easily be generalized to the vector case by considering element-wise maximum operations.

Algorithm 1 Proposed approach

```
for all  $i \in \mathcal{V}', j \in \mathcal{N}_i$, do
    Randomly initialize $z_{i|j}^{(0)} \sim \mathcal{N}(\mu_z, \sigma_z^2)$      ▷ Initialization
    $\text{Node}_j \leftarrow \text{Node}_i(z_{i|j}^{(0)})$
for $t = 0, 1, ...$ do
    for all $i \in \mathcal{V}$  do
        $x_i^{(t+1)} = \frac{-1 - \sum_{j\in\mathcal{N}_i} A_{ij} z_{i|j}^{(t)} + (z_{i|i'}^{(t)} + \frac{1}{2}cs_i)}{c(d_i+1)}$      (4)
        $\forall j \in \mathcal{N}_i : z_{j|i}^{(t+1)} = \frac{1}{2} z_{j|i}^{(t)} + \frac{1}{2}(z_{i|j}^{(t)} + 2cA_{ij}x_i^{(t+1)})$      (5)
        $\text{Node}_{j\in\mathcal{N}_i} \leftarrow \text{Node}_i(x_i^{(t+1)})$      ▷ Broadcast
        for all  $j \in \mathcal{N}_i$ do
            $z_{i|j}^{(t+1)}$ from (5)
        $y_{i|i'}^{(t)} = z_{i|i'}^{(t)} - 2cx_i^{(t+1)} + cs_i$; $y_{i'|i}^{(t)} = z_{i'|i}^{(t)} + cs_i$
        if $y_{i|i'}^{(t)} + y_{i'|i}^{(t)} > 0$ then      ▷ Dummy updates
            $z_{i|i'}^{(t+1)} = \frac{1}{2} z_{i|i'}^{(t)} + \frac{1}{2} y_{i'|i}^{(t)}$      (6)
            $z_{i'|i}^{(t+1)} = \frac{1}{2} z_{i'|i}^{(t)} + \frac{1}{2} y_{i|i'}^{(t)}$      (7)
        else
            $z_{i|i'}^{(t+1)} = \frac{1}{2} z_{i|i'}^{(t)} - \frac{1}{2} y_{i|i'}^{(t)}$      (8)
            $z_{i'|i}^{(t+1)} = \frac{1}{2} z_{i'|i}^{(t)} - \frac{1}{2} y_{i'|i}^{(t)}$      (9)
```

distributed optimization with equality constraints [19]. The main idea is to properly initialize the auxiliary variable $\boldsymbol{z}^{(0)}$, thereby safeguarding the private data from being exposed without sacrificing the output accuracy. Details of the proposed approach are summarized in Alg. 1. Note that the updates at each node $i \in \mathcal{V}'$ can be done in parallel and that no direct collaboration is required between nodes during the computation of these updates, leading to an attractive (parallel) algorithm for optimization in practical networks.

We now analyze the performances of the proposed approach.

A. Output accuracy

When subspace perturbation is applied to inequality-constrained problems, it is shown in [26, Proposition 1] that the optimization variable in A/PDMM, under both equality and inequality constraints, converge to the optimal solution, regardless of the initial values of the auxiliary variable. This ensures that the accuracy of the output is not compromised by the initialization choice of the auxiliary variable. Therefore, our primary focus will be on proving the privacy guarantees.

B. Individual privacy

Given that the eavesdropping adversary holds the information transmitted over all channels given by $\{x_i^{(t+1)} : t \geq 0, i \in \mathcal{V}\} \cup \{z_{i|j}^{(0)} : (i,j) \in \mathcal{E}\}$, and corrupt nodes hold local updates information $\{s_j, z_{j|i}^{(t)}, z_{i|j}^{(t)} : t \geq 0, j \in \mathcal{V}_c, (i,j) \in \mathcal{E}'\}$. Let $\mathcal{T} = \{0, 1, ..., t_{\max}\}$. Given $i \in \mathcal{V}_h$, the individual privacy of honest node i is defined as how much information about the private data s_i can be inferred given the adversaries' knowledge. This is measured by

$$I(S_i;\mathcal{O}) = I(S_i; \{S_j\}_{j\in\mathcal{V}_c}, \{X_j^{(t+1)}\}_{j\in\mathcal{V},t\in\mathcal{T}}, \{Z_{j|k}^{(0)}\}_{(j,k)\in\mathcal{E}}, \{Z_{j|k}^{(t)}, Z_{k|j}^{(t)}\}_{j\in\mathcal{V}_c,(j,k)\in\mathcal{E}',t\in\mathcal{T}}) \tag{10}$$

Without loss of generality, assuming the private data s_is are drawn from independent distributions, our main result is given in Theorem 1, which states that the proposed approach can guarantee (asymptotically) perfect individual privacy even though all other nodes are corrupt, i.e., no information about its private data s_i can be inferred by the passive and eavesdropping adversaries.

Theorem 1. *Given $i \in \mathcal{V}_h$. If*

$$\forall t \in \mathcal{T}: \; z_{i|i'}^{(t)} + z_{i'|i}^{(t)} - 2cx_i^{(t+1)} + 2cs_i \leq 0, \tag{11}$$

then $\lim_{\sigma_z\to\infty} I(S_i;\mathcal{O}) = \lim_{\sigma_z\to\infty} I(S_i; Z_{i|i'}^{(0)} + \frac{1}{2}cS_i) \to 0.$ (12)

Proof. See Appendix A. □

Note that (11) is equivalent to the condition $y_{i|i'}^{(t)} + y_{i'|i}^{(t)} \leq 0$, see (3). In other words, in order to preserve privacy we should avoid data exchange between dummy and regular nodes.

Several remarks are in place here. First, from (11), it is clear that privacy is guaranteed by the honest node's dummy nodes, meaning no honest neighbor is required for privacy assurance. Second, the node with maximum value will not satisfy condition (11) which is to be expected as we require perfect output accuracy so that the value $s_{\max}$ will be eventually available to all nodes. However, for the remaining nodes, the condition for privacy can be satisfied. In the following section, we will demonstrate that it is possible to meet this condition for all iterations by adjusting the convergence parameter c.

IV. Simulation results

Experimental setting: We compare our method with existing privacy-preserving maximum consensus approaches [28], [29]. We generate a random geometric graph (RGG) [37] with $n = 10$ nodes. The private data, i.e. s_i for $i \in \mathcal{V}$, are randomly drawn from a standard normal distribution $\mathcal{N}(0,1)$. The auxiliary variables $z_{i|i'}^{(0)}$ and $z_{i'|i}^{(0)}$ are drawn from $\mathcal{N}(\mu, \sigma^2)$ and $\mathcal{N}(-\mu, \sigma^2)$, respectively. Here μ can take a large value (we use $\mu = 1000$ in the experiments) to ensure that condition (11) is satisfied; a larger value of $z_{i|i'}$ will result in a correspondingly larger value of x_i. To counterbalance the influence of $z_{i|i'}$ in (11), however, we introduce a similar or larger negative value for $z_{i'|i}$ at initialization.

A. Information leakage via mutual information

To visualize information loss in (12) as a function of the variance of the inserted noise, we used NPEET toolbox [38] to estimate the normalized mutual information $I(S_i; Z_{i|i'}^{(0)} + \frac{1}{2}cS_i)/I(S_i;S_i)$ across σ, as depicted Fig. 2. As expected, the information loss decreases notably as σ increases.

B. Performance comparison

We first show that condition (11), required in Theorem 1, can be satisfied at all times by adjusting the convergence parameter c. Fig. 3 shows the LHS of (11) as a function of the

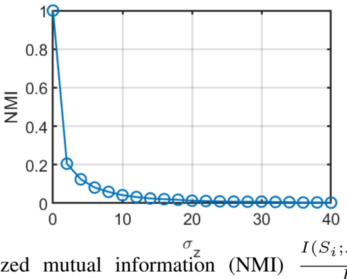

Fig. 2. Normalized mutual information (NMI) $\frac{I(S_i;Z_{i|i'}^{(0)}+\frac{1}{2}cS_i)}{I(S_i;S_i)}$ as a function of variance σ .

iteration number for three different choices of the parameter c. We can see that 1) the blue curve, representing $x_i^{(t)}$ of the node having the maximum value, does not meet condition (11). This is expected as the maximum value will eventually be known to all nodes. 2) For other nodes, a larger parameter c helps to satisfy condition (11), thereby guaranteeing privacy.

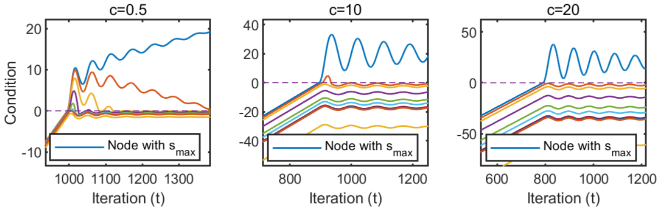

Fig. 3. LHS of (11) as a function of t for three values of c, where the blue lines are the results for the node having the maximum value and the others for nodes having $s_i < s_{\max}$.

In Fig. 4 we compare our proposed approach with two existing algorithms [28], [29] using three privacy levels with different noise variance $\sigma = 10^{-2}, 10^{-1}, 10^{0}$, respectively. It is evident that as the noise increases, [28] exhibits a pronounced deterioration in accuracy while the convergence speed of [29] is affected, highlighting the trade-off between privacy and accuracy. In contrast, our proposed method converges to the optimal result regardless of the noise variance, demonstrating that it does not compromise accuracy for privacy. This is further detailed in Fig. 5 where the convergences of minimum, median and maximum nodes are illustrated.

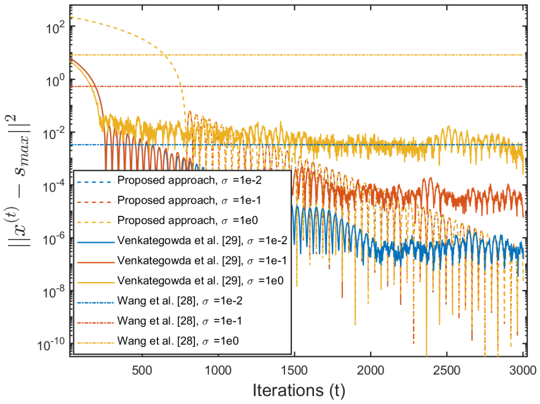

Fig. 4. Performance comparison of the proposed approach with two existing approaches under various privacy levels.

V. Conclusion

In this paper, we proposed a novel privacy-preserving distributed maximum consensus algorithm. Our method involves adding dummy nodes to form an augmented graph and applying inequality constraint-based subspace perturbation, ensuring the privacy of honest participants. Using information-theoretical measures via mutual information, we demonstrated that the proposed approach can guarantee perfect privacy against both eavesdropping and passive adversaries. Furthermore, the method preserves privacy without compromising accuracy, maintaining superior performance. Experimental results further consolidate the superiority of our approach compared to existing methods.

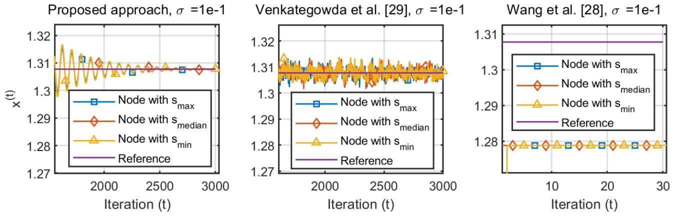

Fig. 5. Convergence of the optimization variable $x^{(t)}$ for three nodes with minimum, median and maximum value of three algorithms, respectively.

Appendix

A. Proof of Theorem 1

Replacing $z_{j|i}^{(t)}$ and $z_{i|j}^{(t)}$ in (5) we obtain

$$z_{j|i}^{(t+1)} - z_{j|i}^{(t)} = cA_{ij}x_i^{(t+1)} - \frac{1}{2}cA_{ij}x_i^{(t)} + \frac{1}{2}cA_{ji}x_j^{(t)}. \quad (13)$$

Moreover, considering the difference $x_j^{(t+2)} - x_j^{(t+1)}$ using (4) and combining with (13) we obtain

$$\begin{aligned} & x_j^{(t+2)} - x_j^{(t+1)} \\ &= \frac{c\sum_{k\in\mathcal{N}_j}(x_k^{(t+1)} - \frac{1}{2}x_k^{(t)} - \frac{1}{2}x_j^{(t)}) + (z_{j|j'}^{(t+1)} - z_{j|j'}^{(t)})}{c(d_j+1)}. \end{aligned} \quad (14)$$

With this, (10) becomes

$$\begin{aligned} & I(S_i;\mathcal{O}) \\ & \overset{(a)}{=} I(S_i;\{S_j, Z_{j|j'}^{(0)}, Z_{j'|j}^{(0)}\}_{j\in\mathcal{V}_c}, \{X_j^{(t+1)}\}_{j\in\mathcal{V},t\in\mathcal{T}}, \{Z_{j|k}^{(0)}\}_{(j,k)\in\mathcal{E}}) \\ & \overset{(b)}{=} I(S_i;\{S_j, Z_{j|j'}^{(0)}, Z_{j'|j}^{(0)}\}_{j\in\mathcal{V}_c}, \{Z_{j|j'}^{(t+2)} - Z_{j|j'}^{(t+1)}\}_{j\in\mathcal{V},t\in\mathcal{T}}, \\ & \quad \{X_j^{(1)}, X_j^{(2)}\}_{j\in\mathcal{V}}, \{Z_{j|k}^{(0)}\}_{(j,k)\in\mathcal{E}}) \\ & \overset{(c)}{=} I(S_i;\{S_j, Z_{j|j'}^{(0)}, Z_{j'|j}^{(0)}\}_{j\in\mathcal{V}_c}, \{X_j^{(1)}, X_j^{(2)}\}_{j\in\mathcal{V}}, \{Z_{j|k}^{(0)}\}_{(j,k)\in\mathcal{E}}) \\ & \overset{(d)}{=} I(S_i;\{S_j, Z_{j|j'}^{(0)}, Z_{j'|j}^{(0)}\}_{j\in\mathcal{V}_c}, \{Z_{j|k}^{(0)}\}_{(j,k)\in\mathcal{E}}, \{Z_{j|j'}^{(0)} + \frac{1}{2}cS_j\}_{j\in\mathcal{V}_h}) \\ & \overset{(e)}{=} I(S_i; Z_{i|i'}^{(0)} + \frac{1}{2}cS_i) \end{aligned}$$

where (a) follows from (5), (6) and (8) since the z variables can be derived from previous z, x_i, and s_i values, (b) follows from (14) since $\{x_j^{(1)}, x_j^{(2)}\}_{j\in\mathcal{V}}$ and $\{z_{j|j'}^{(t+2)} - z_{j|j'}^{(t+1)}\}_{j\in\mathcal{V},t\in\mathcal{T}}$ are sufficient to compute all $\{x_j^{(t+1))}\}_{j\in\mathcal{V},t\in\mathcal{T}}$, and vise versa, (c) assumes that condition (11) is satisfied so that $z_{j|j'}^{(t+2)} - z_{j|j'}^{(t+1)} = c(x_i^{(t+2)} - x_i^{(t+1)})$ can be recursively computed from $\{x_j^{(1)}, x_j^{(2)}\}_{j\in\mathcal{V}}$, (d) holds since $\{x_j^{(1)}, x_j^{(2)}\}_{j\in\mathcal{V}_c}$ can be computed from $z^{(0)}$ and s from corrupt nodes, while (e) holds because $S_j, j \neq i$, and all zs are independent of S_i. Hence, when $\sigma_{Z_{i|i'}} \to \infty$, $Z_{j|j'}^{(0)} + \frac{1}{2}cS_j$ becomes independent of S_i, and thus $I(S_i;\mathcal{O}) \to 0$, thereby completing the proof.

References

[1] G. França and J. Bento, "Distributed optimization, averaging via admm, and network topology," *Proceed. IEEE*, vol. 108, no. 11, pp. 1939–1952, 2020.

[2] R. Zhang and J. Kwok, "Asynchronous distributed admm for consensus optimization," in *Int. Conf. Mach. Learn.* PMLR, pp. 1701–1709, 2014.

[3] D. Deplano, N. Bastianello, M. Franceschelli, and K. H. Johansson, "A unified approach to solve the dynamic consensus on the average, maximum, and median values with linear convergence," in *2023 62nd IEEE Conf. Decis. Control (CDC)*, pp. 6442–6448, 2023.

[4] Z. Huang, S. Mitra, and N. Vaidya, "Differentially private distributed optimization.," in *Proc. Int. Conf. Distrib. Comput. Netw*, pp. 1–10, 2015.

[5] E. Nozari, P. Tallapragada, and J. Cortés, "Differentially private distributed convex optimization via functional perturbation," *IEEE Trans. Control Netw. Syst.*, vol. 5, no. 1, pp. 395–408, 2018.

[6] T. Zhang and Q. Zhu, "Dynamic differential privacy for ADMM-based distributed classification learning," *IEEE Trans. Inf. Forensics Secur.*, vol. 12, no. 1, pp. 172–187, 2016.

[7] X. Zhang, M. M. Khalili, and M. Liu, "Improving the privacy and accuracy of ADMM-based distributed algorithms," *Proc. Int. Conf. Mach. Learn.*, pp. 5796–5805, 2018.

[8] Y. Xiong, J. Xu, K. You, J. Liu and L. Wu, "Privacy preserving distributed online optimization over unbalanced digraphs via subgradient rescaling," *IEEE Trans. Control Netw. Syst.*, 2020.

[9] N. Gupta, J. Katz, N. Chopra, "Privacy in distributed average consensus," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 9515–9520, 2017.

[10] Q. Li, I. Cascudo, and M. G. Christensen, "Privacy-preserving distributed average consensus based on additive secret sharing," in *Proc. Eur. Signal Process. Conf.*, pp. 1–5, 2019.

[11] K. Tjell and R. Wisniewski, "Privacy preservation in distributed optimization via dual decomposition and ADMM," in *Proc. IEEE 58th Conf. Decis. Control.*, pp. 7203–7208, 2020.

[12] K. Tjell, I. Cascudo and R. Wisniewski, "Privacy preserving recursive least squares solutions," in *Proc. Eur. Control Conf.*, pp. 3490–3495, 2019.

[13] Z. Xu and Q. Zhu, "Secure and resilient control design for cloud enabled networked control systems," in *Proc. 1st ACM Workshop Cyber-Phys. Syst.-Secur. Privacy.*, pp. 31–42, 2015.

[14] Q. Li and M. G. Christensen, "A privacy-preserving asynchronous averaging algorithm based on shamir's secret sharing," in *Proc. Eur. Signal Process. Conf.*, pp. 1–5, 2019.

[15] Y. Shoukry et al., "Privacy-aware quadratic optimization using partially homomorphic encryption," in *IEEE 55th Conf. Decis. Control.*, pp. 5053–5058, 2016.

[16] C. Zhang, M. Ahmad, and Y. Wang, "ADMM based privacy-preserving decentralized optimization," *IEEE Trans. Inf. Forensics Secur.*, vol. 14, no. 3, pp. 565–580, 2019.

[17] Q. Li, R. Heusdens and M. G. Christensen, "Convex optimisation-based privacy-preserving distributed average consensus in wireless sensor networks," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, pp. 5895–5899, 2020.

[18] Q. Li, R. Heusdens and M. G. Christensen, "Convex optimization-based privacy-preserving distributed least squares via subspace perturbation," in *Proc. Eur. Signal Process. Conf.*, 2020.

[19] Q. Li, R. Heusdens, and M. G. Christensen, "Privacy-preserving distributed optimization via subspace perturbation: A general framework," *IEEE Trans. Signal Process.*, vol. 68, pp. 5983–5996, 2020.

[20] S. O. Jordan, Q. Li, and R. Heusdens, "Privacy-preserving distributed optimisation using stochastic PDMM," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, pp. 8571–8575, 2024.

[21] Q. Li, R. Heusdens, and M. G. Christensen, "Communication efficient privacy-preserving distributed optimization using adaptive differential quantization," *Signal Process.*, vol. 194, pp. 108456, 2022.

[22] Q. Li, J. S. Gundersen, M. Lopuhaä-Zwakenberg, and R. Heusdens, "Adaptive differentially quantized subspace perturbation (ADQSP): A unified framework for privacy-preserving distributed average consensus," *IEEE Trans. Inf. Forensics Security.*, 2023.

[23] R. Cramer, I. B. Damgård, and J. B. Nielsen, *Secure Multiparty Computation and Secret Sharing*, Cambridge University Press, 2015.

[24] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein, et al., "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, 2011.

[25] G. Zhang and R. Heusdens, "Distributed optimization using the primal-dual method of multipliers," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 4, no. 1, pp. 173–187, 2017.

[26] R. Heusdens and G. Zhang, "Distributed optimisation with linear equality and inequality constraints using pdmm," *IEEE Trans. Signal Inf. Process. Netw.*, 2024.

[27] K. Pillutla, S M Kakade, and Z. Harchaoui, "Robust aggregation for federated learning," *IEEE Trans. Signal Process.*, vol. 70, pp. 1142–1154, 2022.

[28] X. Wang, J. He, P. Cheng, and J. Chen, "Differentially private maximum consensus: Design, analysis and impossibility result," *IEEE Trans. Netw. Sci. Eng.*, vol. 6, no. 4, pp. 928–939, 2018.

[29] N. KD Venkategowda and S. Werner, "Privacy-preserving distributed maximum consensus," *IEEE Signal Process. Lett.*, vol. 27, pp. 1839–1843, 2020.

[30] C. Gratton, N. KD Venkategowda, R. Arablouei, and S. Werner, "Privacy-preserved distributed learning with zeroth-order optimization," *IEEE Trans. Inf. Forensics Secur.*, vol. 17, pp. 265–279, 2021.

[31] P. Cuff and L. Yu, "Differential privacy as a mutual information constraint," in *Proc. 23rd ACM SIGSAC Conf. Comput. Commun. Secur.*, pp. 43–54, 2016.

[32] Q. Li, J. S. Gundersen, R. Heusdens and M. G. Christensen, "Privacy-preserving distributed processing: Metrics, bounds, and algorithms," in *IEEE Trans. Inf. Forensics Secur.*, pp. 2090–2103, 2021.

[33] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, "Local privacy and statistical minimax rates," in *Proc. IEEE Annu. Symp. Found. Comput. Sci.*, pp. 429–438, 2013.

[34] P. Kairouz, S. Oh, and P. Viswanath, "Extremal mechanisms for local differential privacy," in *Adv. Neural Inf. Process. Syst.*, pp. 2879–2887, 2014.

[35] Q. Li, J. S. Gundersen, K. Tjell, R. Wisniewski, and M. G. Christensen, "Privacy-preserving distributed expectation maximization for gaussian mixture model using subspace perturbation," in *IEEE Proc. Int. Conf. Acoust., Speech, Signal Process.*, pp. 4263–4267, 2022.

[36] T. M. Cover and J. A. Tomas, *Elements of information theory*, John Wiley & Sons, 2012.

[37] M. Penrose, *Random geometric graphs*, vol. 5, OUP Oxford, 2003.

[38] G. V. Steeg, "Npeet," https://github.com/gregversteeg/NPEET.