

Two for the price of one

communication efficient and privacy-preserving distributed average consensus using quantization

Li, Qiongxiu; Lopuhaä-Zwakenberg, Milan; Heusdens, Richard; Christensen, Mads Græsbøll

Publication date

2022

Document Version

Final published version

Published in

30th European Signal Processing Conference, EUSIPCO 2022 - Proceedings

Citation (APA)

Li, Q., Lopuhaä-Zwakenberg, M., Heusdens, R., & Christensen, M. G. (2022). Two for the price of one: communication efficient and privacy-preserving distributed average consensus using quantization. In *30th European Signal Processing Conference, EUSIPCO 2022 - Proceedings* (pp. 2166-2170). (European Signal Processing Conference; Vol. 2022-August). European Signal Processing Conference, EUSIPCO.

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Two for the price of one: communication efficient and privacy-preserving distributed average consensus using quantization

Qiongxiu Li^{1,*}, Milan Lopuhaä-Zwakenberg^{2,*}, Richard Heusdens³, Mads Græsbøll Christensen¹

¹ Audio Analysis Lab, CREATE, Aalborg University, Denmark, qili,mgc@create.aau.dk

²University of Twente, The Netherlands, m.a.lopuhaa@utwente.nl

³ Netherlands Defence Academy and Delft University of Technology, The Netherlands, r.heusdens@tudelft.nl

Abstract—Both communication overhead and privacy are main concerns in designing distributed computing algorithms. It is very challenging to address them simultaneously as encryption methods required for privacy-preservation often incur high communication costs. In this paper, we argue that there is a fundamental link between communication efficiency and privacy-preservation through quantization. Based on the observation that quantization, which can save communication bandwidth, will introduce error into the system, we propose a novel privacy-preserving distributed average consensus algorithm which uses the error introduced by quantization as noise to obfuscate the private data for protecting it from being revealed to others. Similar to existing differential privacy based approaches, the proposed approach is robust and has low computational complexity in dealing with two widely considered adversary models: the passive and eavesdropping adversaries. In addition, the method is generally applicable to many distributed optimizers, like PDMM and (generalized) ADMM. We conduct numerical simulations to validate that the proposed approach has superior performance compared to existing algorithms in terms of accuracy, communication bandwidth and privacy.

Index Terms—Distributed average consensus, privacy, wireless sensor networks, communication, ADMM, PDMM

I. INTRODUCTION

With the emergence of advanced microprocessor design and wireless communication technologies, there is a huge growth in distributed computing systems. The distributed average consensus problem has served as a fundamental building block for different fields such as optimization [1], robotics, signal processing and machine learning, where the goal is to compute the global average over all participants' data in such distributed systems. Due to the fact that the involved participants' data often comes from personal devices such as mobile phones and tablets [2], [3], these data carries sensitive personal information thus imposes privacy concerns. In addition, the involved computing devices are usually limited in communication bandwidth and computing resources. Overall, it is crucial to design novel distributed average consensus solutions that are both privacy-preserving and lightweight in terms of computation and communication cost.

* The first two authors contributed equally to this work. This research has been partially funded by NWO grant 628.001.026, the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 101008233, and ERC consolidator grant 864075 "Caesar".

Existing algorithms only address the above-mentioned challenges partially. The first type of approach [4], [5] addresses the privacy issue by adopting the homomorphic encryption [6] to ensure computation over an encrypted domain. However, homomorphic encryption incurs a high cost on both communication and computational expenses [7]. The other type of approach, referred to as the information-theoretical or noise insertion approach, is more computationally lightweight. The main idea is to design a lightweight encryption function by inserting noise to mask the private data before sharing to others. Consider a security attack where there are untrustworthy participants in the system, an algorithm is very robust in privacy if it is able to protect privacy even though there are many untrustworthy participants. For each participant, the extreme trustless scenario is that all other participants are not trustworthy. It has been shown that under such an extreme trustless scenario, the two desired performances: output accuracy and guaranteed privacy, cannot be achieved simultaneously in the context of distributed average consensus [8], [9]. Therefore, different algorithms have to prioritize one over another. Based on the prioritized performance, information-theoretical approaches can be further classified into two types. Differentially private algorithms [9]–[11], which apply the general concept of differential privacy [12], achieve privacy guarantee under the extreme trustless scenario by compromising accuracy. Alternatively, other algorithms [13]–[21] achieve output accuracy but they require the additional assumption that the number of untrustworthy participants should not exceed a certain threshold. The main idea is to design the noise insertion process in particular ways such that the accuracy of average output is not affected, e.g., all inserted noise are coordinated to sum up to zero [13]–[18], or inserting noise into a specific subspace that has no effects on accuracy [20], [21].

The last challenge, i.e., communication efficiency, is often overlooked in existing noise insertion approaches by assuming infinite precision. However, we remark that there is often a fundamental trade-off between communication bandwidth and privacy level for noise insertion approaches. A higher privacy level usually requires a larger amount of noise, resulting in an increase in noise entropy and thus the amount of bits. In this paper, we propose a novel approach which is able to address communication efficiency and privacy-preservation

simultaneously. A typical way to save communication cost is to apply quantization schemes. We observe that quantization will introduce additional error to the system; while the main idea of information-theoretical approaches is to insert noise for privacy preservation. Therefore, we propose to link them by adopting the error introduced by quantization as noise to protect privacy. To the best of our knowledge, this is the first distributed average consensus approach which makes use of quantization to ensure both information-theoretical privacy and communication efficiency.

II. PRELIMINARIES AND PROBLEM DEFINITION

In this section, we review necessary fundamentals and concepts to define the problem to be solved.

A. Distributed average consensus over networks

Let the undirected graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ model a distributed network where $\mathcal{N} = \{1, 2, \dots, n\}$ denotes node set and \mathcal{E} denotes edge set and $m = |\mathcal{E}|$. Moreover, let $\mathcal{N}_i = \{j \mid \{i, j\} \in \mathcal{E}\}$ denote the set of neighboring nodes and $d_i = |\mathcal{N}_i|$. Assume each node i has local data s_i . For simplicity, we assume that s_i is scalar-valued but the results hold for arbitrary dimensions. The goal of distributed average consensus is to compute the global average over the network:

$$s_{\text{ave}} = n^{-1} \sum_{i \in \mathcal{N}} s_i, \quad (1)$$

in a distributed manner, i.e., each node is only allowed to communicate with its neighboring nodes.

B. Private data definition and adversary models

We define the local data s_i of each node as the private data to be protected. The adversary model is an important concept when analyzing privacy. It specifies what kinds of security attack the algorithm is addressing. Besides the well-known eavesdropping adversary which attacks the system by eavesdropping the communication channels between nodes, another commonly considered model in distributed networks is the passive or honest-but-curious adversary. It attacks the system by colluding a number of nodes in the network. These nodes are called corrupt nodes and are assumed to follow the algorithm instructions but collect information together. With all information collected by the corrupt nodes, the passive adversary aims to infer the private data of the rest non-corrupt, referred to as honest nodes. Note that the extreme trustless scenario described in the introduction means that there are $n - 1$ corrupt nodes when considering the passive adversary.

C. Main requirements

We note that there are three main requirements. 1) **Output correctness**: at the end of the algorithm, each node i would like to obtain the average result s_{ave} . 2) **Individual privacy**: throughout the algorithm, each node's private data s_i , should be protected from being revealed to both eavesdropping and passive adversaries. 3) **Communication cost**: the algorithm should have low communication cost.

D. Local perturbation method and privacy metrics

An important technique to ensure privacy is *local perturbation*, in which each node i privately generates a random number r_i , and engages in the distributed average consensus protocol with $\tilde{s}_i = s_i + r_i$. Since the adversaries only learn about s_i through \tilde{s}_i , it provides privacy even if all other $n - 1$ nodes are corrupt. Such privacy guarantee comes at an accuracy cost, as the protocol will now compute $\tilde{s}_{\text{ave}} = s_{\text{ave}} + r_{\text{ave}}$, with r_{ave} being the average value of the inserted noise.

Many metrics exist to measure the privacy leakage. A commonly used one is Local Differential Privacy (LDP) [22], which gives strong *plausible deniability* privacy guarantees. However, because LDP gives a requirement on all possible inputs, it often requires additional assumption on the private data and can be difficult to realize in practice. Furthermore, it comes at a large accuracy cost. In this paper we use mutual information [23] as a more practically applicable metric to quantify the information leakage. The mutual information $I(S_i; \tilde{S}_i)$ measures how much information regarding S_i is revealed by knowing \tilde{S}_i and vice versa, where we used uppercase letters to denote random variables and lowercase letters to indicate the corresponding realization. Mutual information can be seen as a relaxed version of LDP [24].

III. DISTRIBUTED OPTIMIZER AND QUANTIZATION

In this section we will briefly introduce the fundamentals for the proposed approach.

A. Distributed optimizer

The goal of distributed optimizers is to solve decomposable optimization problems over a network in a distributed manner. Many distributed optimizers have been proposed such as ADMM [25] and PDMM [26]. It has been shown that ADMM and PDMM are closely related using monotone operator theory and operator splitting techniques [26] (see [27] for details on monotone operator theory). For both methods, the update equations at iteration $t = 0, 1, \dots$ are given by

$$x_i^{(t+1)} = \arg \min_{x_i} (f_i(x_i) + \sum_{j \in \mathcal{N}_i} z_{i|j}^{(t)} B_{i|j} x_i + \frac{c d_i}{2} x_i^2) \quad (2)$$

$$z_{j|i}^{(t+1)} = \theta z_{j|i}^{(t)} + (1 - \theta)(z_{i|j}^{(t)} + 2c B_{i|j} x_i^{(t+1)}), \quad (3)$$

where c denotes the penalty parameter for controlling the convergence rate, $x_i^{(t)}$ denotes the primal (optimization) variable of node i at iteration t , $f_i(x_i)$ denotes the local objective function of node i , assumed to be convex. For each edge $e_l = (i, j) \in \mathcal{E}$, there are two so-called auxiliary variables $z_l = z_{i|j}$ and $z_{l+m} = z_{j|i}$. $B \in \mathbb{R}^{m \times n}$ is related to the graph incidence matrix: $B_{i|j} = 1$, $B_{j|i} = -1$ if and only if $(i, j) \in \mathcal{E}$ and $i < j$. The constant $\theta \in [0, 1]$ is used for controlling the averaging weight of Peaceman-Rachford splitting where $\theta = 0.5$ corresponds to ADMM, and $\theta = 0$ to PDMM. In the following we will use PDMM as an example to explain the main idea but the conclusions hold for all $\theta \in [0, 1]$.

B. Adaptive differential quantization

Adaptive differential quantization was introduced in [28], [29] to perform distributed optimization at a low communication cost without compromising the optimization accuracy. It is motivated by the fact that for fixed point iterations the difference of successive updates will converge to zero, thus the entropy of the difference will decrease as the number of iteration increases. Therefore, it first defines an adaptive decreasing cell-width for quantization:

$$\Delta^{(t)} = \gamma^t \Delta^{(0)}, \quad (4)$$

where $\Delta^{(0)}$ denotes the initialized cell-width, $\gamma \in (0, 1)$ is a constant for controlling the decreasing rate of cell-width. Let l denote the word length of the quantized message. With cell-width $\Delta^{(t)}$, define a l -bit uniform (mid-rise) quantization function $Q^{(t)}: \mathbb{R} \rightarrow \Delta^{(t)}(a + 1/2)$, where $a \in \mathcal{A} = \{-2^{l-1}, -2^{l-1} + 1, \dots, 2^{l-1} - 1\}$. Each input is mapped into its nearest representation value among all $\{\Delta^{(t)}(a + 1/2)\}_{a \in \mathcal{A}}$.

The defined adaptive quantizer is then used to quantize a difference variable of successive iterations defined as

$$\forall (i, j) \in \mathcal{E}: v_{i|j}^{(t)} = \begin{cases} z_{i|j}^{(0)} & \text{if } t = 0 \\ z_{i|j}^{(t)} - z_{i|j}^{(t-1)} & \text{if } t \geq 1. \end{cases} \quad (5)$$

Denote \hat{v} as the quantized v , we have

$$\hat{v}_{i|j}^{(t)} = Q^{(t)}\left(z_{i|j}^{(t)} - z_{i|j}^{(t-1)} + \delta_{i|j}^{(t)}\right), \quad (6)$$

where $\delta_{i|j}^{(t)}$ denotes the so-called subtractive dithering noise which is added to the input before quantization at the transmission side and later be subtracted from the quantized output in the receiving side. More specifically, node j first transmits $\hat{v}_{i|j}^{(t)}$ to node i . Node i then subtracts the dithering noise, i.e., $\hat{v}_{i|j}^{(t)} = \hat{v}_{i|j}^{(t)} - \delta_{i|j}^{(t)}$ and updates $\hat{z}_{i|j}^{(t)}$ locally by

$$\hat{z}_{i|j}^{(t)} = \begin{cases} \hat{v}_{i|j}^{(0)} & \text{if } t = 0 \\ \hat{z}_{i|j}^{(t-1)} + \hat{v}_{i|j}^{(t)} & \text{if } t \geq 1. \end{cases} \quad (7)$$

After constructing $\hat{z}_{i|j}^{(t)}$, node i then uses it to replace the unquantized $z_{i|j}^{(t)}$ in (2) and (3) to update $x_i^{(t+1)}$ and $\{z_{j|i}^{(t+1)}\}_{j \in \mathcal{N}_i}$ locally. The usage of dithering noise $\delta_{i|j}^{(t)}$, which is made uniformly distributed on $[-\frac{\Delta^{(t)}}{2}, \frac{\Delta^{(t)}}{2}]$, is to ensure that the quantization error

$$n_{i|j}^{(t)} = \hat{z}_{i|j}^{(t)} - z_{i|j}^{(t)} \quad (8)$$

is independent of $v_{i|j}^{(t+1)}$, and thus of $z_{i|j}^{(t+1)}$. This has the advantage of resulting better accuracy [30]. The subtractive dithering noise can be implemented by sharing seeds of a pseudo-random generator before starting the algorithm.

IV. PROPOSED APPROACH

We first formulate the distributed average consensus problem as the following optimization function:

$$\begin{aligned} \min_{x_i} \quad & \sum_{i \in \mathcal{N}} \frac{1}{2} \|x_i - s_i\|_2^2 \\ \text{s.t.} \quad & \forall (i, j) \in \mathcal{E}: x_i = x_j, \end{aligned} \quad (9)$$

where the optimum solution is given by $\forall i \in \mathcal{N}: x_i^* = s_{\text{ave}}$. Using PDMM and adaptive differential quantization, the updating functions are given by

$$x_i^{(t+1)} = \frac{s_i - \sum_{j \in \mathcal{N}_i} B_{i|j} \hat{z}_{i|j}^{(t)}}{1 + cd_i} \quad (10)$$

$$\forall j \in \mathcal{N}_i: z_{j|i}^{(t+1)} = \hat{z}_{i|j}^{(t)} + 2cB_{i|j}x_i^{(t+1)}. \quad (11)$$

One thing to note here is that we slightly revise the adaptive quantization scheme by setting a minimum cell-width Δ_{\min} . That is, instead of using (4), we have

$$\Delta^{(t)} = \max\{\gamma^t \Delta^{(0)}, \Delta_{\min}\}. \quad (12)$$

This is to guarantee a minimum amount of quantization noise for privacy under the extreme trustless scenario, i.e., $n - 1$ corrupt nodes when dealing with the passive adversary (we will give detailed proof in the following subsection). Details of the proposed approach are summarized in Algorithm 1.

Algorithm 1 Proposed approach

Quantization parameters: $\Delta^{(0)}, \Delta_{\min}, \gamma, l$

Optimization output : $x_i^{(t_{\max})}$ for every node $i \in \mathcal{N}$

For each node $i \in \mathcal{N}$:

while $0 \leq t < t_{\max}$ **do**

$\Delta^{(t)} \leftarrow (12), \{\hat{v}_{j|i}^{(t)}\}_{j \in \mathcal{N}_i} \leftarrow (6);$

Transmit $\hat{v}_{j|i}^{(t)}$ to each neighbor $j \in \mathcal{N}_i$;

Receive $\hat{v}_{i|j}^{(t)}$ from all neighbors $j \in \mathcal{N}_i$;

$\{\hat{z}_{i|j}^{(t)}\}_{j \in \mathcal{N}_i} \leftarrow (7), x_i^{(t+1)} \leftarrow (10), \{z_{j|i}^{(t+1)}\}_{j \in \mathcal{N}_i} \leftarrow (11);$

end

A. Performance analysis

By inspecting Algorithm 1, we can see that the only information needed to be transmitted along the network is the quantized $\{\hat{v}_{i|j}^{(t)}\}_{(i,j) \in \mathcal{E}, t \geq 0}$. Based on (7) we have

$$\hat{z}_{i|j}^{(t)} = \sum_{\tau=0}^{t-1} \hat{v}_{i|j}^{(\tau)}. \quad (13)$$

To quantify the individual privacy of node i , i.e., the amount of information about its private data s_i is learned by the adversaries, we must first inspect what information is available to the adversaries. Denote \mathcal{V}_e and \mathcal{V}_p as the set of information available to the eavesdropping and passive adversary, respectively. We first consider the eavesdropping adversary. A typical way to address this adversary is to securely encrypt the communication channels such that the transmitted information cannot be eavesdropped. However, such channel encryption methods [31] are often communication and computationally expensive. Since our goal is to develop lightweight solutions, we assume that there is no channel encryption, i.e., the eavesdropping adversary is able to eavesdrop all transmitted information: $\mathcal{V}_e = \{\hat{v}_{j|k}^{(t)}\}_{(j,k) \in \mathcal{E}, t \geq 0}$. Regarding the passive adversary, as we mentioned before that different algorithms have to choose between accurate average result and privacy guarantee against $n - 1$ corrupt nodes

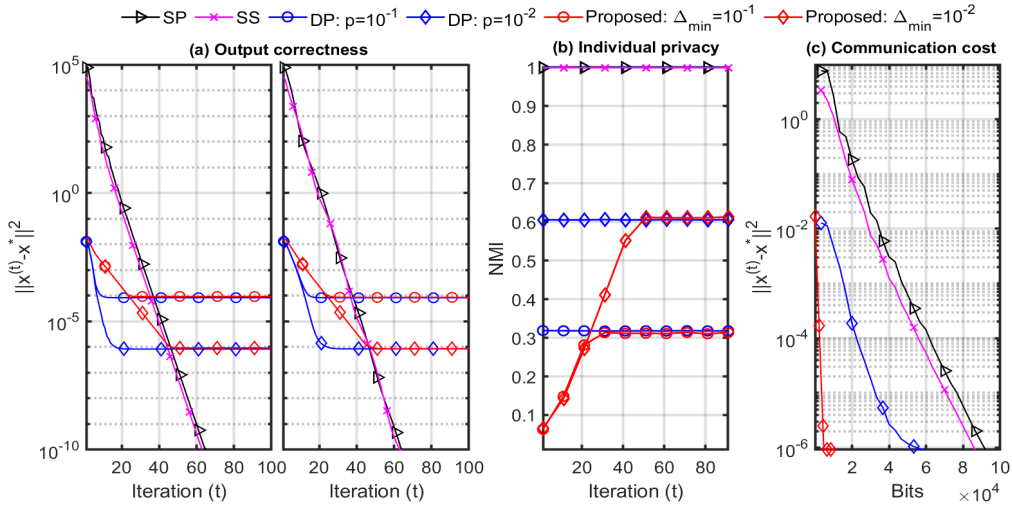


Fig. 1: Performance comparisons of existing SP, SS, DP based approaches and the proposed one when dealing with $n - 1$ corrupt nodes, where for the latter two we consider two different privacy levels. (a) Output correctness: mean square error (MSE) of optimization variable as a function of iteration numbers based on PDMM (left) and ADMM (right), respectively. (b) Individual privacy: normalized mutual information (NMI) as a function of iteration numbers. (c) Communication cost: MSE as a function of the amount of bits used for transmitting all messages.

since they cannot be achieved simultaneously. In the proposed approach, we prioritize the latter by considering $n - 1$ corrupt nodes. Denote node i as the only honest node. The set of information available to the passive adversary is given by:

$$\mathcal{V}_p = \{s_j, x_j^{(t)}\}_{j \in \mathcal{N} \setminus \{i\}, t \geq 1} \cup \{\hat{v}_{j|k}^{(t)}\}_{(j,k) \in \mathcal{E}, t \geq 0}. \quad (14)$$

We can see that $\mathcal{V}_e \subset \mathcal{V}_p$. By inspecting (11) and (8), we conclude that $x_i^{(t+1)} = (\hat{z}_{j|i}^{(t+1)} - n_{j|i}^{(t+1)} - \hat{z}_{i|j}^{(t)}) / (2cB_{i|j})$ and inserting into (10) yields

$$\begin{aligned} \forall j \in \mathcal{N}_i : s_i + \frac{(1+cd_i)n_{j|i}^{(t+1)}}{2cB_{i|j}} \\ = \frac{(1+cd_i)(\hat{z}_{j|i}^{(t+1)} - \hat{z}_{i|j}^{(t)})}{2cB_{i|j}} + \sum_{k \in \mathcal{N}_i} B_{i|k} \hat{z}_{i|k}^{(t)}. \end{aligned} \quad (15)$$

Note that based on (13), all terms in the RHS of equality can be computed using the information in \mathcal{V}_p . We then conclude that regarding the private data s_i of the honest node i , the adversaries observe $\{s_i + \frac{(1+cd_i)n_{j|i}^{(t+1)}}{2cB_{i|j}}\}_{j \in \mathcal{N}_i}$ at each iteration. Denote $\frac{(1+cd_i)}{2c} = c_i$, and $m_{i|j}^{(t+1)} = n_{j|i}^{(t+1)} / B_{i|j}$. Then the individual privacy of node i at iteration $t + 1$ is given by

$$I(S_i; \{S_i + c_i M_{i|j}^{(t+1)}\}_{j \in \mathcal{N}_i}). \quad (16)$$

As the number of iterations increases, the above mutual information will first increase and finally converge when the quantization cell-width reaches its minimum Δ_{\min} . Note that the above individual privacy has a similar form as the local perturbation or relaxed DP approaches using mutual information as a metric, which is given by $I(S_i; \tilde{S}_i) = I(S_i; S_i + R_i)$. Overall, we conclude that the proposed quantization based algorithm is able to guarantee individual privacy under both eavesdropping adversary and passive adversary with $n - 1$

corrupt nodes with a lightweight communication cost. By increasing the minimum cell-width Δ_{\min} , the privacy level is higher, but the average result will be less accurate. We will validate this by numerical results in the coming section.

V. NUMERICAL RESULTS

In this section, we demonstrate numerical results to validate the superior performance of the proposed approach by comparing to three existing approaches including subspace perturbation (SP) [20], [21], secret sharing (SS) [16], [17], (relaxed) differential privacy (DP) [9]–[11] based approaches. The performance is compared in terms of the three requirements specified in Section II-C. As for the metrics, we use the mean square error (MSE) of the optimization variable to quantify the output correctness, normalized mutual information (NMI) for individual privacy, and the amount of bits transmitted over the network $l(2m)t$ for calculating the communication cost.

We simulated a graph with $n = 10$ nodes. All private data are independent and uniformly distributed over $[-0.5, 0.5]$ and we ran 10^4 Monte Carlo simulations. For the proposed approaches, we set $l = 2$ for the quantization, i.e., each transmitted message costs only two bits. Since no quantization is considered in the existing approaches, in the implementation we use the default MATLAB double precision floating-point format, i.e., $l = 64$. The penalty parameter is set as $c = 1$. We consider $n - 1$ corrupt nodes case. For simplicity, we assume that the only honest node has one neighboring node, i.e., $d_i = 1$ such that (16) reduces to $I(S_i; S_i + M_{i|j}^{(t+1)})$ as $c_i = 1$. In order to demonstrate the connection of the proposed approach and relaxed DP approaches as they all provide privacy guarantee against $n - 1$ corrupt nodes, in

the experiments we aim to compare their privacy levels while achieving the same output accuracy. Therefore, we assume that the noise inserted in DP approach is uniformly distributed among interval $[-\frac{p}{2}, \frac{p}{2}]$, and set $p = \Delta_{\min}$.

The results are shown in Fig. 1. As shown in plot (a) the proposed approach is generally applicable to both PDMM (left) and ADMM (right). Due to limited space, in plot (b) and (c) we only show results for PDMM. By inspecting (a) and (b) we can see that when considering $n - 1$ corrupt nodes, the existing SP and SS based approaches are able to achieve accurate average but individual privacy is not protected at all. As expected, the proposed approach ends up with similar accuracy and privacy level as DP based approaches by setting $p = \Delta_{\min}$. In both algorithms, if we increase the amount of noise, i.e., p in the DP approaches and Δ_{\min} in the proposed approach, the privacy level is higher but the average result becomes less accurate. Hence, the proposed approach achieves similar output correctness and individual privacy compared to DP based approaches. Finally, from plot (c) we can see that, as expected, the proposed approach significantly reduces the amount of communication cost compared to all existing approaches. This is because each transmitted message only requires $l = 2$ bits for the proposed quantization based approach, which is much lightweight compared to the $l = 64$ bits for existing approaches.

VI. CONCLUSION

In this paper, we proposed a novel distributed average consensus approach that is both communication efficient and privacy-preserving by using an adaptive differential quantization technique. With the help of quantization, the communication cost can be reduced while the private data can be masked by quantization noise and thus being protected. Numerical results demonstrate that the proposed approach has superior performance compared to three types of existing approaches. In particular, considering the maximum number of corrupt nodes, the proposed approach is able to achieve similar output correctness and individual privacy performance as DP based approaches, but the communication cost is much lower.

REFERENCES

- [1] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *IEEE Proc.*, vol. 95, no. 1, pp. 215-233, 2007.
- [2] M. Anderson, *Technology device ownership, 2015*, Pew Research Center, 2015.
- [3] J. Poushter and others, "Smartphone ownership and internet usage continues to climb in emerging economies," *Pew Research Center*, vol. 22, pp. 1-44, 2016.
- [4] R. C. Hendriks, Z. Erkin, and T. Gerkmann, "Privacy-preserving distributed speech enhancement for wireless sensor networks by processing in the encrypted domain," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, pp. 7005-7009, 2013.
- [5] M. H. Ruan, M. Ahmad, Y. Q. Wang, "Secure and privacy-preserving average consensus," in *Proc. Workshop Cyber-Phys. Syst. Secur. Privacy*, pp. 123-129, 2017.
- [6] P. Paillier, "Public-key cryptosystems based on composite degree residuosity classes," in *EUROCRYPT*, pp. 223-238, 1999.

- [7] R. L. Lagendijk, Z. Erkin, and M. Barni, "Encrypted signal processing for privacy protection: Conveying the utility of homomorphic encryption and multiparty computation," *IEEE Signal Process. Magazine*, vol. 30, no. 1, pp. 82-105, 2013.
- [8] Q. Li, J. S. Gundersen, R. Heusdens and M. G. Christensen, "Privacy-preserving distributed processing: Metrics, bounds, and algorithms," in *IEEE Trans. Inf. Forensics Security*, vol. 16, pp. 2090-2103, 2021.
- [9] E. Nozari, P. Tallapragada, and J. Cortés, "Differentially private average consensus: Obstructions, trade-offs, and optimal algorithm design," *Automatica*, vol. 81, pp. 221-231, 2017.
- [10] M. Kefayati, M. S. Talebi, B. H. Khalajand H. R. Rabiee, "Secure consensus averaging in sensor networks using random offsets," in *Proc. of the IEEE Int. Conf. on Telec., and Malaysia Int. Conf. on Commun.*, pp. 556-560, 2007.
- [11] Z. Huang, S. Mitra, and G. Dullerud, "Differentially private iterative synchronous consensus," in *ACM workshop Privacy electron. Soc.*, pp. 81-90, 2012.
- [12] C. Dwork, F. McSherry, K. Nissim, A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Proc. Theory of Cryptography Conf.*, pp. 265-284, 2006.
- [13] N. E. Manitaru and C. N. Hadjicostis, "Privacy-preserving asymptotic average consensus," in *ECC*, pp. 760-765, 2013.
- [14] Y. Mo and R. M. Murray, "Privacy preserving average consensus," *IEEE Trans. Automat. Contr.*, vol. 62, no. 2, pp. 753-765, 2017.
- [15] P. Braca, R. Lazzaretto, S. Marano, and V. Matta, "Learning with privacy in consensus + obfuscation," *IEEE Signal Process. Lett.*, vol. 23, no. 9, pp. 1174-1178, 2016.
- [16] N. Gupta, J. Katz, N. Chopra, "Privacy in distributed average consensus," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 9515-9520, 2017.
- [17] N. Gupta, J. Kat and N. Chopra, "Statistical privacy in distributed average consensus on bounded real inputs," in *ACC*, pp 1836-1841, 2019.
- [18] Y. Guo and Y. Gong, "Practical collaborative learning for crowdsensing in the internet of things with differential privacy," in *Proc. IEEE Conf. Commun. Netw. Secur. (CNS)*, pp.1-9, 2018.
- [19] 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.
- [20] 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.
- [21] Q. Li, R. Heusdens and M. G. Christensen, "Privacy-preserving distributed optimization via subspace perturbation: A general framework," in *IEEE Trans. Signal Process.*, vol. 68, pp. 5983 - 5996, 2020.
- [22] K. Nissim S. Raskhodnikova S. P. Kasiviswanathan, H. K. Lee and A. Smith, "What can we learn privately?," *SIAM Journal on Computing*, vol. 40, no. 3, pp. 793-826, 2011.
- [23] T. M. Cover and J. A. Tomas, *Elements of information theory*, John Wiley & Sons, 2012.
- [24] G. Barthe and B. Kopf, "Information-theoretic bounds for differentially private mechanisms," in *IEEE 24th Computer Security Foundations Symposium*, 2011, pp. 191-204.
- [25] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein, et al., "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine learning*, vol. 3, no. 1, pp. 1-122, 2011.
- [26] T. Sherson, R. Heusdens, W. B. Kleijn, "Derivation and analysis of the primal-dual method of multipliers based on monotone operator theory," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 5, no. 2, pp 334-347, 2018.
- [27] E. Ryu, S. P. Boyd, "Primer on monotone operator methods," *Appl. Comput. Math.*, vol. 15, no. 1, pp. 3-43., 2016.
- [28] D. H. M. Schellekens, T. Sherson, and R. Heusdens, "Quantisation effects in PDMM: A first study for synchronous distributed averaging," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, pp. 4237-4241, 2017.
- [29] J. A. G. Jonkman, T. Sherson, and R. Heusdens, "Quantisation effects in distributed optimisation," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, pp. 3649-3653, 2018.
- [30] L. Schuchman, "Dither signals and their effect on quantization noise," *IEEE Transactions on Communication Technology*, vol. 12, no. 4, pp. 162-165, 1964.
- [31] D. Dolev, C. Dwork, O. Waarts, M. Yung, "Perfectly secure message transmission," *J. Assoc. Comput. Mach.*, vol. 40, no. 1, pp. 17-47., 1993.