Convex Optimisation-Based Privacy-Preserving Distributed Average Consensus in Wireless Sensor Networks

Li, Qiongxiu ; Heusdens, Richard; Christensen, Mads Græsbøll

DOI

[10.1109/ICASSP40776.2020.9053348](https://doi.org/10.1109/ICASSP40776.2020.9053348)

Publication date

2020

Document Version

Accepted author manuscript

Published in

ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)

Citation (APA)

Li, Q., Heusdens, R., & Christensen, M. G. (2020). Convex Optimisation-Based Privacy-Preserving Distributed Average Consensus in Wireless Sensor Networks. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP): Proceedings* (pp. 5895-5899). IEEE . <https://doi.org/10.1109/ICASSP40776.2020.9053348>

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.

CONVEX OPTIMISATION-BASED PRIVACY-PRESERVING DISTRIBUTED AVERAGE CONSENSUS IN WIRELESS SENSOR NETWORKS

Qiongxu Li*

Richard Heusdens[†]

Mads Græsbøll Christensen*

* Audio Analysis Lab, CREATE, Aalborg University, Denmark, {qili, mgc}@create.aau.dk

[†] Circuits and Systems group, Delft University of Technology, The Netherlands, r.heusdens@tudelft.nl

ABSTRACT

In many applications of wireless sensor networks, it is important that the privacy of the nodes of the network be protected. Therefore, privacy-preserving algorithms have received quite some attention recently. In this paper, we propose a novel convex optimization-based solution to the problem of privacy-preserving distributed average consensus. The proposed method is based on the primal-dual method of multipliers (PDMM), and we show that the introduced dual variables of the PDMM will only converge in a certain subspace determined by the graph topology and will not converge in the orthogonal complement. These properties are exploited to protect the private data from being revealed to others. More specifically, the proposed algorithm is proven to be secure for both passive and eavesdropping adversary models. Finally, the convergence properties and accuracy of the proposed approach are demonstrated by simulations which show that the method is superior to the state-of-the-art.

Index Terms— Distributed average consensus, privacy, wireless sensor networks, convex optimisation, primal-dual method of multipliers

1. INTRODUCTION

Advances in wireless communication technology and embedded microprocessor design have enabled a huge growth of distributed computing systems, including also wireless sensor networks (WSNs). Average consensus, which is an essential building block of such distributed systems, has been intensively investigated for decades, and it has been applied in various fields such as automatic control, signal processing, robotics and optimisation [1]. To solve the average consensus problem in distributed networks, many (iterative) algorithms have been proposed [2–10]. The methods work by iteratively exchanging information between computational units (i.e., nodes/agents), whereby the network eventually reaches a consensus. The data exchange required in these algorithms can lead to privacy problems, as it is becoming clear that there is no real separation between the identity of individuals and their data [11]. Therefore, it is crucial to protect the data held by each node as the private data for being revealed to others.

An algorithm is called secure or privacy-preserving if it is able to protect the private data during the algorithm execution. Existing privacy-preserving distributed average consensus algorithms can be classified into two classes: computationally secure algorithms and information-theoretically secure algorithms. Computational security is defined in terms of computational hardness: secrets cannot be reconstructed efficiently under the condition that so-called malicious adversaries are computationally limited. Computationally secure algorithms [12–16] usually apply techniques from secure multiparty computation [17] such as homomorphic encryption (HE) [18, 19] and garbled circuit (GC) [20, 21], where computations are performed in the encrypted domain. However, these algorithms are computationally demanding and have high a communication bandwidth. This

makes it difficult to apply them in resource constrained applications like WSNs.

In contrast to the aforementioned computationally expensive algorithms, the information-theoretically secure algorithms are quite lightweight by comparison, as they simply insert noise to obfuscate the private data. Moreover, information-theoretic security has a stronger security guarantee than computational security as it is robust against a computationally unlimited adversary. Depending on the amount of information about the private data obtained by the adversary, information-theoretically secure algorithms can be further classified into two classes. The first class contains algorithms using secret sharing, whereby perfect security is achieved [22]. It possesses the strongest security guarantees. No information regarding the private data is revealed as the information obtained by the adversary is statistically independent of the private data. However, it requires prior knowledge about the network. The second class of algorithms achieves a weaker form of security, called ϵ -statistical security, which implies that the information obtained by the adversary is not totally independent of the private data but only results in a slightly better posterior guessing probability than the prior probability. Most ϵ -statistical security algorithms [23–25] adopt differential privacy [26, 27] to obfuscate the private data with independent noise. However, as shown in [25], differential privacy-based approaches cannot obtain the exact average and privacy at the same time. One way to circumvent the trade-off between accuracy and privacy is to guarantee that the inserted noise adds up to zero. Some algorithms [28–30] insert noise having a geometrically decreasing variance over iterations and guarantee that the inserted noise adds up to zero. Some other algorithms [31–33] rely on a trusted third party to obtain the zero-sum property. However, a trusted third party is hard to implement in ad hoc networks including also many WSNs.

As discussed above, the existing information-theoretically secure algorithms have some limitations, such as requiring prior knowledge of the network, the zero-sum property of the inserted noise, or the existence of a trusted third party. To address these limitations, we propose a convex optimisation-based method. To explain the basic concept, we show how it can be applied in the primal-dual method of multipliers (PDMM) [10, 34] which is an iterative algorithm for solving constrained convex optimisation problems. The concept can, however, also be applied to other convex optimisation methods, for example ADMM-based algorithms. As we shall see, the proposed method has a number of attractive properties: 1) the proposed algorithm obtains asymptotically perfect security and requires no trusted party nor prior knowledge about the network; 2) exact consensus and privacy can be obtained simultaneously; 3) the algorithm does not need zero-sum noise insertion but only a proper initialisation of the dual variables; 4) the convergence rate is independent of the privacy level; 5) the algorithm is secure under both passive and eavesdropping adversaries; and 6) the privacy of any honest node is guaranteed as long as it has one honest neighbour.

2. PRELIMINARIES AND PROBLEM DEFINITION

In this section, we will define the problem at hand and introduce some important definitions and concepts.

2.1. Distributed average consensus

Let $G = (V, E)$ denote a simple graph, where $V = \{1, 2, \dots, n\}$ and $E = \{e_1, \dots, e_m\} \subseteq V \times V$ denote the set of nodes and edges, respectively. The neighbourhood of node i is denoted as $N_i = \{j \in V \mid (i, j) \in E\}$ and the degree of node i is denoted by $d_i = |N_i|$. Finally, let $A \in \mathbb{R}^{n \times n}$ denote the adjacency matrix of the graph defined as $A_{ij} = 1$ if and only if $(i, j) \in E$, and let $B \in \mathbb{R}^{m \times n}$ denote the incidence matrix defined as $B_{li} = B_{ij} = 1$ if and only if $e_l = (i, j) \in E$ and $i < j$ and $B_{li} = B_{ij} = -1$ if and only if $e_l = (i, j) \in E$ and $i > j$. Distributed average consensus aims to estimate the average of all the initial state values given by

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

with s_i the initial state value of node i , without any centralised coordination. For simplicity, we will assume that s_i is a scalar but the results can easily be generalised to arbitrary dimensions.

2.2. Privacy concern and adversary model

In this work, the initial state value of each node is the private data to be protected. Most algorithms consider a passive adversary model (also known as the honest-but-curious model) where the instructions of the protocol are followed, but the so-called corrupted nodes might collude and attempt to deduce information about the initial state values of the other honest nodes from the messages they receive. The eavesdropping adversary is usually neglected in existing approaches since eavesdropping can be prevented by using channel encryption [35]. However, channel encryption is computationally expensive. For iterative algorithms where the communication channels between nodes are used many times, channel encryption is, therefore, less attractive. We thus assume that the communication in the network is performed through non-secure channels, except for the communication during the initialisation of the network.

2.3. Problem definition

The goal of privacy-preserving distributed average consensus algorithms is to design a protocol that jointly computes the average of all initial state values while protecting them from being revealed in the process. We thus have the following two requirements which need to be satisfied simultaneously:

- 1) Correctness: at the end of the algorithm, each node has obtained the average result $s_{\text{ave}} = n^{-1} \sum_{i \in V} s_i$.
- 2) Individual privacy: throughout the execution of the algorithm, the initial state value held by each honest node is protected against both passive and eavesdropping adversaries.

Some remarks are in order here. The adversary always knows the sum of the initial state values of the honest nodes, as it can be deduced from the average result and the initial states values of the corrupted nodes. Therefore, revealing this sum is unavoidable [17]. Furthermore, for incomplete (i.e., not fully connected) networks, as in the case in many practical networks, the partial sums of the honest nodes in each (connected) subgraph will be revealed as well, something that is also unavoidable for any information-theoretically private protocol [36, 37].

The corrupted nodes aim to infer the initial state value s_i of node i . Let s_i denote a realisation of a random variable S_i having differential entropy $h(S_i)$, assuming it exists¹, and let $g^{(k)}(S_i)$ denote the

¹In the case that S_i is a discrete random variable, the conditions are given in terms of the Shannon entropy $H(S_i)$.

information sent out at iteration k by node i . We will measure the amount of privacy by

$$I(S_i; g^{(k)}(S_i)) = h(S_i) - h(S_i \mid g^{(k)}(S_i)), \quad (2)$$

where $I(\cdot; \cdot)$ denotes mutual information [38]. Note that $I(S_i; g^{(k)}(S_i)) = 0$ corresponds to perfect security in the sense that $h(S_i \mid g^{(k)}(S_i)) = h(S_i)$ so that S_i and $g^{(k)}(S_i)$ are statistically independent, while $I(S_i; g^{(k)}(S_i)) < \epsilon$, where $\epsilon > 0$, corresponds to ϵ -statistical security. Again, having perfect security at every iteration does not necessarily imply that $I(S_i; g^{(k)}(S_i), \dots, g^{(0)}(S_i)) = 0$ since in the end the adversary is able to compute partial sums of connected subgraphs, but nothing else beyond that.

3. PRIMAL-DUAL METHOD OF MULTIPLIERS

The proposed approach is based on the primal-dual method of multipliers (PDMM), an instance of Peaceman-Rachford splitting of the extended dual problem (see [34] for details). PDMM can, like ADMM, be used for iteratively solving constrained convex optimisation problems. The PDMM update equations are given by

$$\begin{aligned} x^{(k+1)} &= \arg \min_x \left(f(x) + \lambda^{(k)T} PCx + \frac{c}{2} \|Cx + PCx^{(k)}\|_2^2 \right), \\ \lambda^{(k+1)} &= P\lambda^{(k)} + c(Cx^{(k+1)} + PCx^{(k)}), \end{aligned} \quad (3)$$

where k denotes the iteration index, $x^{(k)} \in \mathbb{R}^n$ is the primal variable, $\lambda^{(k)} \in \mathbb{R}^{2m}$ the dual variable, $f(x)$ the objective function to be minimised, $C \in \mathbb{R}^{2m \times n}$ a matrix related to the graph's incidence matrix B , and $P \in \mathbb{R}^{2m \times 2m}$ a symmetric permutation matrix exchanging the first m with the last m rows. The $c > 0$ is a constant controlling the convergence rate. The vector λ contains the dual variables controlling the constraints; for each edge $(i, j) \in E$ there are two node variables $\lambda_{i|j}$ and $\lambda_{j|i}$, one for each node i and j , respectively, where $\lambda(l) = \lambda_{i|j}$ and $C_{li} = B_{ij}$ if and only if $e_l = (i, j) \in E$ and $i < j$, and $\lambda(l+m) = \lambda_{j|i}$, $C_{(l+m)i} = B_{ij}$ if and only if $e_l = (i, j) \in E$ and $i > j$. Note that $C + PC = [B^T \ B^T]^T$ and $\forall (i, j) \in E : \lambda_{j|i} = (P\lambda)_{i|j}$.

Consider the update of two successive λ -updates, given by

$$\lambda^{(k+2)} = \lambda^{(k)} + c(Cx^{(k+2)} + 2PCx^{(k+1)} + Cx^{(k)}), \quad (4)$$

since $P^2 = I$. Let $H = \text{ran}(C) + \text{ran}(PC)$ where $\text{ran}(\cdot)$ denotes the range, and let Π_H denote the orthogonal projection onto H . By inspection of (4), we conclude that every two PDMM updates only affect $\Pi_H \lambda \in H$ and leave $(I - \Pi_H)\lambda \in H^\perp$, $H^\perp = \text{null}(C^T) \cap \text{null}((PC)^T)$ unchanged, where $\text{null}(\cdot)$ denotes the null space. Moreover, by inspecting (3), we conclude that the x -update is independent of $(I - \Pi_H)\lambda$ since $\lambda^T(I - \Pi_H)PC = 0$. As a consequence, the component $(I - \Pi_H)\lambda$ will only be permuted every iteration and therefore not converge. We will refer to $\Pi_H \lambda$ and $(I - \Pi_H)\lambda$ as the converging and non-converging component of the dual variable, respectively.

4. PROPOSED APPROACH

The distributed average consensus problem can be formulated as an optimisation problem where we minimise the objective function

$$f(x) = \frac{1}{2} \|x - s\|_2^2, \quad (5)$$

where $s = (s_1, \dots, s_n)^T$, subject to the constraint that $x_i = x_j$ for all $(i, j) \in E$. The solution is given by $x^* = s_{\text{ave}}(1, \dots, 1)^T$. That is, all nodes in the network eventually know the average. The PDMM update equation (3) for this problem is then given by

$$x^{(k+1)} = (I + cD)^{-1} \left(s + cAx^{(k)} - C^T P\lambda^{(k)} \right), \quad (6)$$

where $D = C^T C$ is the degree matrix of the underlying graph and $C^T P C = -A$. The update equations for node i then become

$$x_i^{(k+1)} = \frac{s_i + \sum_{j \in N_i} (c x_j^{(k)} - B_{i|j} \lambda_{j|i}^{(k)})}{1 + c d_i}, \quad (7)$$

$$\forall j \in N_i : \lambda_{i|j}^{(k+1)} = \lambda_{j|i}^{(k)} + c B_{i|j} (x_i^{(k+1)} - x_j^{(k)}). \quad (8)$$

From (8) we can see that the update of the dual variables only depends on $\lambda_{j|i}^{(k)}$, $x_j^{(k)}$ and $x_i^{(k+1)}$, of which $\lambda_{j|i}^{(k)}$ and $x_j^{(k)}$ are already available at node j . Therefore, after broadcasting $x_i^{(k+1)}$, all neighbouring nodes can construct $\lambda_{i|j}^{(k+1)}$ and the dual variables do not need to be transmitted at all, except for the initialisation, as all $\lambda_{j|i}^{(0)}$ s need to be known at the first iteration.

As mentioned before, the non-converging component $(I - \Pi_H) \lambda^{(k)}$ will only be permuted every iteration so that

$$\lambda^{(k)} \rightarrow \lambda^* + \begin{cases} (I - \Pi_H) \lambda^{(0)}, & k \text{ even,} \\ P(I - \Pi_H) \lambda^{(0)}, & k \text{ odd,} \end{cases} \quad (9)$$

where λ^* is given by

$$\lambda^* = - \left(\begin{matrix} C^T \\ (PC)^T \end{matrix} \right)^\dagger \left(\begin{matrix} \nabla f(x^*) + c C^T C x^* \\ \nabla f(x^*) + c C^T P C x^* \end{matrix} \right) + c C x^*, \quad (10)$$

where $(\cdot)^\dagger$ denotes the Moore-Penrose pseudo inverse. As a consequence, if we initialise the dual variable λ in such a way that the non-converging component $(I - \Pi_H) \lambda^{(0)}$ sufficiently obfuscates the initial state value, the primal variables will converge to s_{ave} while the initial state value itself cannot be inferred, assuming there is at least one honest neighbour. We will prove this claim more formally in what follows.

4.1. Correctness

As shown in [34], the primal variable $x^{(k)}$ will converge geometrically to x^* for arbitrary initialisation $x^{(0)}$ and $\lambda^{(0)}$, thereby proving the correctness of the algorithm.

4.2. Individual privacy

We will now proceed to prove that the proposed algorithm protects the individual privacy under both passive and eavesdropping adversaries. As we can see, each node transmits only the primal variable $x_i^{(k+1)}$ to all of its neighbours and does not reveal its initial state value s_i directly. To analyse the privacy properties of the proposed algorithm, let V_c and V_h denote the set of corrupted and honest nodes, respectively. With this, the numerator of (7) can be expressed as

$$\begin{aligned} s_i + \sum_{j \in N_i} (c x_j^{(k)} - B_{i|j} \lambda_{j|i}^{(k)}) &= \\ s_i + \sum_{j \in N_i} c x_j^{(k)} - \sum_{j \in N_i \cap V_h} B_{i|j} \lambda_{j|i}^{(k)} - \sum_{j \in N_i \cap V_c} B_{i|j} \lambda_{j|i}^{(k)}. \end{aligned} \quad (11)$$

At convergence, x^* is known and λ^* can be calculated through (10). Hence, by inspection of (9) and (11), we conclude that the adversary can infer about the initial state value s_i from observing $x_i^{(k+1)}$ is the term given by

$$s_i - \sum_{j \in N_i \cap V_h} B_{i|j} \left(P^k (I - \Pi_H) \lambda^{(0)} \right)_{j|i}, \quad (12)$$

and we conclude that, as long as $N_i \cap V_h \neq \emptyset$, we can obfuscate the initial state value by introducing uncertainty in $(I - \Pi_H) \lambda^{(0)}$.

Algorithm 1 Privacy-preserving PDMM

- 1: Each node $i \in V$ initialises its primal and dual variables. The dual variables are initialised with random numbers having sufficiently large variance (depending on the required privacy level), whereas the primal variables can be initialised arbitrarily.
  - 2: Each node $i \in V$ communicates its dual variables $\lambda_{i|j}^{(0)}$ to its neighbour $j \in N_i$ through secure channels [35].
  - 3: **while** $\|x^{(k)} - x^*\|_2 < \text{threshold}$ **do**
  - 4:   Activate a node uniformly at random, say node i , updates its primal variable $x_i^{(k+1)}$ according to (7).
  - 5:   Node i broadcasts $x_i^{(k+1)}$ to all of its neighbours $j \in N_i$ (through non-secure channels).
  - 6:   After receiving $x_i^{(k+1)}$ by the neighbours, the dual variables $\lambda_{i|j}^{(k+1)}$ are updated using (8).
  - 7: **end while**
-

To quantitatively measure the amount of information carried by $x_i^{(k)}$ about s_i , consider both $x_i^{(k)}$ and s_i as realisations of the random variables $X_i^{(k)}$ and S_i , respectively. We will analyse the mutual information $I(S_i; X_i^{(k)})$ between S_i and $X_i^{(k)}$ for which we need the following result.

Proposition 1. *Let X and Y be independent continuous random variables with $\text{var}(X), \text{var}(Y) < \infty$ and let $Z = X + Y$. Then*

$$\lim_{\text{var}(Y) \rightarrow \infty} I(X; Z) = 0,$$

assuming $I(X; Z)$ exists.

Proof. Let $\gamma = 1/(\text{var}(Y))^{\frac{1}{2}}$ and define $Y' = \gamma Y$. Hence, Y' has unit variance. Since mutual information is invariant under scaling, we have $I(X; Z) = I(X; X + Y) = I(\gamma X; \gamma X + Y')$. As a consequence, we have

$$\begin{aligned} \lim_{\text{var}(Y) \rightarrow \infty} I(X; Z) &= \lim_{\gamma \rightarrow 0} I(\gamma X; \gamma X + Y') \\ &= I(0; Y') = 0. \end{aligned} \quad \square$$

By applying Proposition 1, we can conclude that the mutual information $I(S_i; X_i^{(k)})$ can be made arbitrarily small by increasing the variance of the random variable representing the λ -contribution in (12). That is, let $\lambda^{(0)}$ be a realisation of the random variable $\Lambda^{(0)}$. Then we have $I(S_i; X_i^{(k)}) = 0$ if

$$\exists j \in N_i \cap V_h : \text{var} \left(((I - \Pi_H) \Lambda^{(0)})_{j|i} \right) \rightarrow \infty. \quad (13)$$

Hence, the proposed algorithm obtains asymptotically perfect security. A summary of the complete privacy-preserving PDMM algorithm is given in Algorithm 1.

Some remarks are in order here. Firstly, since the dual variables are not transmitted at all, except during initialisation for which we need secure communication, no encryption is needed during the execution of the algorithm. Secondly, a necessary condition for achieving privacy is that $N_i \cap V_h \neq \emptyset$. That is, node i requires at least one honest neighbour. In the case the graph is complete, this means that the algorithm is secure up to $n - 2$ malicious nodes in the network. Thirdly, although we proved that the mutual information is zero under the condition of (13), the variance of the dual variables cannot be made infinitely large. Therefore, information about the initial state variables will be leaked upon receiving the primal variables. To have an indication of the amount of leakage in practice,

Fig. 1: Convergence of the primal variable, the converging component and non-converging component of the dual variable in PDMM with two different initialisations.

consider the following example of two independent Gaussian distributed random variables X and Y and their sum $Z = X + Y$. The differential entropy of a Gaussian random variable with variance σ^2 is given by $\frac{1}{2} \log(2\pi e \sigma^2)$, so that $I(X; Z) = h(Z) - h(Y) = \frac{1}{2} \log(1 + \sigma_X^2/\sigma_Y^2)$. Hence, if we have $\sigma_Y^2/\sigma_X^2 = 100$ (the range of Y is approximately 10 times the range of X), the information leakage is only 0.007 bits. Fourthly, in order to satisfy (13), a necessary condition is that $\lambda^{(0)} \cap H^\perp \neq \emptyset$. By inspection of the matrix C , we conclude that the matrix $[C, PC] \in \mathbb{R}^{2m \times 2n}$ can be considered as the incidence matrix of a bipartite graph having $2n$ nodes. As a consequence, we have that $\text{rank}([C, PC]) \leq 2n - 1$ and we conclude that $\dim(H) \leq 2n - 1$ and thus $H^\perp \neq \emptyset$. Hence, if we randomly initialise $\lambda^{(0)}$, we have $(I - \Pi_H)\lambda^{(0)} \neq 0$ with probability one. Last but not least, the proposed algorithm can also be applied to other convex optimisation methods such as ADMM and related algorithms where the update equations have a similar structure.

5. EXPERIMENTAL RESULTS

Now we proceed to evaluate the performance of the proposed algorithm by simulations in terms of the mean square error (MSE) of primal and dual variables as a function of transmission number. We generated a random geometric network with $n = 10$ nodes where two nodes can communicate if their distance is within a radius r satisfying $r^2 = 2 \frac{\log n}{n}$, thereby guaranteeing a connected graph with probability at least $1 - \frac{1}{n^2}$ [39]. For simplicity, we use uniform distribution as an example to demonstrate the results, where the initial state values s_i are uniformly distributed in the interval $[0, 1]$.

Figure 1 shows the convergence behavior of PDMM for different initialisations. The red lines show the proposed PDMM algorithm in which $x^{(0)}$ is initialised with all zeros and $\lambda^{(0)}$ is randomly initialised with uniformly distributed noise in the interval $[0, 100]$. The green lines show results where the dual variable is initialised in H such that $\lambda^{(0)} \cap H^\perp = \emptyset$, which implies that the initial state values are not protected. The star, square, and circle marker show the convergence of $x^{(k)}$, $(I - \Pi_H)\lambda^{(k)}$ and $\Pi_H\lambda^{(k)}$, respectively. We see that for both initialisations $x^{(k)}$ and $\Pi_H\lambda^{(k)}$ converge to the optimal solutions x^* and λ^* , respectively. The magnitude of $(I - \Pi_H)\lambda^{(k)}$, on the other hand, does not converge. As a consequence, the proposed algorithm protects the initial state value by obfuscating it with a high-variance non-converging component

Fig. 2: Convergence of the proposed PDMM and state-of-the-art algorithms under three different noise levels.

$(I - \Pi_H)\lambda^{(k)}$. Note that the green line with square marker is not visible since $(I - \Pi_H)\lambda^{(k)} = 0$ for all k .

Figure 2 shows a comparison of the proposed PDMM approach with popular state-of-the-art information-theoretically secure algorithms including differential privacy (DP) [25] and the correlated noise insertion approach (CNI) [30], where we compare the effect of adding noise on the convergence rate of the algorithm. We considered three different noise levels: $\Gamma = 0, 10^2$, and 10^4 , where Γ denotes the ratio of noise variance to the variance of initial state value. The case $\Gamma = 0$ corresponds to the situation where no noise is added so that the initial state values are not protected. In the other cases we inserted noise having an initial range approximately 10 and 100 times the range of initial state values, therefore we have $\Gamma = 10^2$ and 10^4 , respectively. We observe, as expected, that the accuracy of the differential privacy approach (black lines) decreases with increasing noise variance and that for $\Gamma \neq 0$ the algorithm does not converge anymore. That is, with differential privacy, there is a trade-off between privacy and accuracy. As for correlated noise insertion (blue lines), high accuracy is obtained in the end (the algorithm is guaranteed to converge) but the convergence rate slows down with increasing noise variance. The convergence rate of the proposed PDMM-based algorithm (red lines), on the other hand, is independent of the noise level since the convergence rate of PDMM depends on the graph topology and not on the initialisation; increasing the noise variance will only result in a higher initial error.

6. CONCLUSIONS

In this paper, we proposed a novel lightweight privacy-preserving distributed average consensus algorithm for WSNs based on PDMM, a convex optimisation algorithm. By simply initialising the dual variable with random numbers, the non-converging component of the dual variable will obfuscate the initial state values, thereby protecting them from being revealed. We showed that the proposed algorithm achieves asymptotically perfect security under a passive adversary, where the privacy is guaranteed as long as there is at least one honest neighbour. For an eavesdropping adversary, the proposed algorithm does not require secure channel encryption in the network except for the initialisation step. Compared to existing information-theoretically secure algorithms, the proposed algorithm has no trade-off between accuracy and privacy, and converges at a rate independent of the amount of inserted noise and, thus, of the level of privacy.

7. 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] L. Xiao, S. Boyd, "Fast linear iterations for distributed averaging," *Syst. Control Lett.*, vol. 53, no. 1, pp. 65–78, 2004.
- [3] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized gossip algorithms," *IEEE Trans. Inf. Theory*, vol. 52, no. 6, pp. 2508–2530, 2006.
- [4] A. G. Dimakis, S. Kar, J. M. Moura, M. G. Rabbat, and A. Scaglione, "Gossip algorithms for distributed signal processing," *Proc. IEEE*, vol. 98, no. 11, pp. 1847–1864, 2010.
- [5] C.M. Bishop, *Pattern recognition and machine learning*, Springer, 2006.
- [6] J. Pearl, *Reverend Bayes on inference engines: A distributed hierarchical approach*, Proc. 1982 Am. Assoc. Artificial Intell., pp. 133–136, 1982.
- [7] S. Aliaksei and K. Soummya and M. José MF, "Finite-time distributed consensus through graph filters," in *ICASSP*, pp. 1080–1084, 2014.
- [8] S. Santiago and M. Antonio G and R. Alejandro, "Optimal graph-filter design and applications to distributed linear network operators," *IEEE Trans. Signal Process.*, vol. 65, no. 15, pp. 4117–4131, 2017.
- [9] 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.
- [10] G. Zhang and R. Heusdens, "Distributed optimization using the primal-dual method of multipliers," *IEEE Trans. Signal Process.*, vol. 4, no. 1, pp. 173–187, 2018.
- [11] . D. Sarwate and K. Chaudhuri, "Signal processing and machine learning with differential privacy: Algorithms and challenges for continuous data," *IEEE Signal Process. Magazine*, vol. 30, no. 5, pp. 86–94, 2013.
- [12] R. C. Hendriks, Z. Erkin, and T. Gerkmann, "Privacy preserving distributed beamforming based on homomorphic encryption," in *EUSIPCO*, pp. 1–5, 2013.
- [13] R. C. Hendriks, Z. Erkin, and T. Gerkmann, "Privacy-preserving distributed speech enhancement for wireless sensor networks by processing in the encrypted domain," in *ICASSP*, pp. 7005–7009, 2013.
- [14] 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.
- [15] C. Zhang, M. Ahmad, and Y. Wang, "ADMM based privacy-preserving decentralized optimization," *IEEE Trans. Inf. Forensics Security*, vol. 14, no. 3, pp. 565–580, 2019.
- [16] F. Hanzely, J. Konečný, N. Loizou, P. Richtárik, and D. Grishchenko, "Privacy preserving randomized gossip algorithms," *arXiv preprint arXiv:1706.07636*, 2017.
- [17] R. Cramer, I. B. Damgrd, and J. B. Nielsen, *Secure Multiparty Computation and Secret Sharing*, Cambridge University Press, 2015.
- [18] P. Paillier, "Public-key cryptosystems based on composite degree residuosity classes," in *EUROCRYPT*, pp. 223–238, 1999.
- [19] I. Damgård, V. Pastro, N. Smart, and S. Zakarias, "Multiparty computation from somewhat homomorphic encryption," in *Advances in Cryptology—CRYPTO*, pp. 643–662. Springer, 2012.
- [20] A. C. Yao, "Protocols for secure computations," in *FOCS*, pp. 160–164, 1982.
- [21] A. C. Yao, "How to generate and exchange secrets," in *FOCS*, pp. 162–167, 1986.
- [22] Q. Li, I. Cascudo, and M. G. Christensen, "Privacy-preserving distributed average consensus based on additive secret sharing," in *EUSIPCO*, to appear, 2019.
- [23] 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.
- [24] Z. Huang, S. Mitra, and G. Dullerud, "Differentially private iterative synchronous consensus," in *ACM workshop Privacy electron. Soc.*, pp. 81–90, 2012.
- [25] 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.
- [26] C. Dwork, "Differential privacy," in *ICALP*, pp. 1–12, 2006.
- [27] 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.
- [28] N. E. Manitaras and C. N. Hadjicostis, "Privacy-preserving asymptotic average consensus," in *ECC*, pp. 760–765, 2013.
- [29] Y. Mo and R. M. Murray, "Privacy preserving average consensus," *IEEE Trans. Automat. Contr.*, vol. 62, no. 2, pp. 753–765, 2017.
- [30] J. He, L. Cai, C. Zhao, P. Cheng, X. Guan, "Privacy-preserving average consensus: privacy analysis and algorithm design," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 5, no. 1, pp. 127–138, 2019.
- [31] P. Braca, R. Lazzeretti, S. Marano, and V. Matta, "Learning with privacy in consensus + obfuscation," *IEEE signal process. Lett.*, vol. 23, no. 9, pp. 1174–1178, 2016.
- [32] M. T. Hale, M. Egerstedt, "Differentially private cloud-based multi-agent optimization with constraints," in *Proc. American Control Conf.*, pp. 1235–1240, 2015.
- [33] M. T. Hale, M. Egerstedt, "Cloud-enabled differentially private multiagent optimization with constraints," *IEEE Trans. Control Netw. Syst.*, vol. 5, no. 4, pp. 1693–1706, 2018.
- [34] 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.*, 2018.
- [35] D. Dolev, C. Dwork, O. Waarts, M. Yung, "Perfectly secure message transmission," *J. Assoc. Comput. Mach.*, vol. 40, no. 1, pp. 17–47, 1993.
- [36] G. Kreitz, M. Dam, and D. Wikstrom, "Practical private information aggregation in large networks," in *In: Aura, T., Järvinen, K., Nyberg, K. (eds.) NordSec. LNCS*, vol. 7127, pp. 89–103, 2010.
- [37] Beimel, A., "On private computation in incomplete networks," *Distrib. Comput.* 19(3), 237–252, 2007.
- [38] T.M. Cover and J.A. Tomas, *Elements of information theory*, John Wiley & Sons, 2012.
- [39] J. Dall and M. Christensen, "Random geometric graphs," *Physical review E*, vol. 66, no. 1, pp. 016121, 2002.