

A Privacy-Preserving Asynchronous Averaging Algorithm based on State Decomposition

Metin Calis¹, Richard Heusdens^{1,2}, Richard C. Hendriks¹

¹*Delft University of Technology, Department of Microelectronics, Delft, Netherlands*

²*Netherlands Defence Academy*

{m.calis,r.heusdens,r.c.hendriks}@tudelft.nl

Abstract—Average consensus algorithms are used in many distributed systems such as distributed optimization, sensor fusion and the control of dynamic systems. Consensus algorithms converge through an explicit exchange of state variables. In some cases, however, the state variables are confidential. In this paper, a privacy-preserving asynchronous distributed average consensus method is proposed, which decomposes the initial values into two states; alpha states and beta states. These states are initialized such that their sum is twice the initial value. The alpha states are used to communicate with the other nodes, while the beta states are used internally. Although beta states are not shared, they are used in the update of the alpha states. Unlike differential privacy based methods, the proposed algorithm achieves the exact average consensus, while providing privacy to the initial values. Compared to the synchronous state decomposition algorithm, the convergence rate is improved without any privacy compromise. As the variances of coupling weights become infinitely large, the semi-honest adversary does not have any range to estimate the initial value of the nodes given that there is at least one coupling weight hidden from the adversary.

Index Terms—Privacy-preserving averaging, Distributed averaging, State decomposition, Asynchronous averaging

I. INTRODUCTION

Consensus problems in dynamic systems have been a topic of interest that have found usage in many research areas allowing multiple agents to reach an agreement through local information exchange between the agent and its neighbors [1]. Some of these research areas are sensor fusion [2], control of swarms and flocks [3] [4] and alignment problems [5]. The traditional consensus algorithms explicitly exchange their state variables to solve a common function. However, for some consensus problems such as the multi- rendezvous problem [6] or energy management in smart grids [7], the initial states can be confidential. In the former, the agents might not want to reveal their initial locations, while in the latter, the energy companies might not want to reveal their individual generation rates. The challenge to solve the consensus problem while giving individual nodes a privacy guarantee initiated the fairly new privacy-preserving distributed optimization research area.

II. RELATED WORK

The research directed towards solving the consensus problem while preserving the privacy of initial values can be categorized into two approaches: cryptographic [8] [9] and non-cryptographic methods [10] [11] [12] [13]. Most of the cryptographic methods use homomorphic encryption [14] to

encrypt the states that are being transferred. In control and real-time dynamic systems where processing time is limited, cryptographic methods become infeasible due to the time the encryption and the decryption takes. To reduce the time and complexity, privacy-preserving non-cryptographic consensus methods have been proposed.

Differential privacy based approaches, e.g. [10] [15], trade accuracy for privacy. Nodes add noise to the transmitted states and provide a differential privacy guarantee as defined in [16] or in [17] for continuous data observations. However, as proven by [10], differential privacy and exact consensus cannot be achieved simultaneously. To achieve exact consensus, the adversary should be prevented from eavesdropping on all of the communication of the targeted node. Noise-obfuscation methods [11] [13] use a topological constraint on the graph and achieve exact consensus through the addition of correlated noise to the transmitted states. As the added noise is zero-sum and decaying in magnitude over iterations, the exact average can be achieved. Due to this noise insertion mechanism, noise has to be added to the transmitted states until convergence. In [18] it is shown that achieving the exact average without any privacy compromise is possible by perturbing the states once, before starting the consensus process. In this paper, an asynchronous privacy-preserving consensus average method is proposed that extends the state decomposition approach [12] by showing that the same convergence rate as the standard distributed averaging methods [19] can be achieved, while promising identical privacy guarantee.

Privacy-preserving asynchronous averaging has been investigated in [18] [20] and [21]. In [18] Shamir's secret sharing scheme [22] is used to hide the initial values from the adversaries. Perfect secrecy is provided in clique-based graphs as long as the selected clique at each iteration includes at least two honest nodes. In [20] an additive secret sharing scheme is proposed as a preprocessing step where the mutual information between the processed values and the initial values approaches zero with the increasing noise variance. The added noise is arranged such that the exact consensus can be achieved. The algorithm can be applied to any connected graph given that there are at least two honest nodes. In [21] three noise-obfuscation based methods are proposed that add correlated and decaying noise to the transmitted states at each iteration. The authors analyze the rate and conditions of convergence without quantifying the provided privacy. The

proposed method in this work, differs from [18] and [20] in the sense that the network does not include any channel encryptions and the added noise is additive and multiplicative. A combination of additive and multiplicative noise might affect the provided privacy when the noise is sampled from distributions with small variances. On the other hand, the proposed method differs from [21] as the noise is added to the system for a bounded amount of time and the provided privacy is quantified using an information-theoretic analysis.

III. PRELIMINARIES

We represent an undirected graph G as $G = (V, E)$ with its node set $V = \{v_1, v_2, \dots, v_N\}$ and its edge set $E \subset V \times V$. The i th component of the vector $x[k] = [x_1[k], x_2[k], \dots, x_N[k]]$ represents the state of node v_i at iteration k . The set of neighbors of node v_i is $N_i = \{v_j \in V : (v_i, v_j) \in E\}$ and its cardinality is given by $|N_i|$. The goal is to compute $x_{\text{ave}} = \frac{1}{N} \sum_{j=1}^N x_j[0]$ using an asynchronous algorithm, while hiding the initial states. In this work, we analyze the privacy against attacks by a passive adversary and an eavesdropper. Throughout the paper the following assumption is made.

Assumption 1: The graph is connected, undirected and there are no channel encryptions in the network.

One way to solve the average consensus problem in an asynchronous fashion is to use the randomized gossip algorithm [19] using the iterations

$$x_i[k+1] = x_i[k] + \frac{1}{2}(x_j[k] - x_i[k]). \quad (1)$$

Under Assumption 1 and using independent doubly stochastic weight matrices [23], it is proven that the state variables converge almost surely to $x_{\text{ave}} = \frac{1}{N} \sum_{j=1}^N x_j[0]$.

A semi-honest adversary is defined to be a node in the network who follows the protocol steps correctly but tries to gain more information by collecting the data they receive. An eavesdropper on the other hand, taps arbitrarily any communication channel. However, the eavesdropper has no access to internal state variables that are not shared in the system.

IV. ASYNCHRONOUS STATE DECOMPOSITION

Each node decomposes its state value, say $x_i[0] \in \mathbb{R}$ into two substates $x_i^\alpha[0] \in \mathbb{R}$ and $x_i^\beta[0] = 2x_i[0] - x_i^\alpha[0]$ resulting in an increase in the number of nodes from N to $2N$. Notice that these additional nodes are virtual nodes. Substate $x^\alpha[k]$ is used in the interaction with the other nodes, while $x^\beta[k]$ is used as an internal update. Although $x^\beta[k]$ is never shared, it is used in the evaluation of $x^\alpha[k]$. Using the state decomposition approach [12], the update equations become,

$$\begin{aligned} x_i^\alpha[k+1] &= x_i^\alpha[k] + \frac{1}{3}(x_j^\alpha[k] - x_i^\alpha[k]) + \\ &\quad \frac{1}{3}(x_i^\beta[k] - x_i^\alpha[k]), \\ x_i^\beta[k+1] &= x_i^\beta[k] + \frac{1}{3}(x_i^\alpha[k] - x_i^\beta[k]), \end{aligned} \quad (2)$$

subject to $x_i^\alpha[0] + x_i^\beta[0] = 2x_i[0]$.

Fixing the coupling weights to be $1/3$ limits the privacy that can be provided. For this reason, two phases are proposed: initialization phase and the consensus phase. In the initialization phase, the coupling weights are selected from the set of real numbers under the condition that the sum of all state variables does not change. Selecting the coupling weights from the set of real numbers introduces randomness to the system that will protect the initial values. As the sum of the state variables does not change, the exact consensus can still be achieved.

During the consensus phase, the update equations are identical to (1). As privacy is already established in the initialization phase, the motivation is to let the nodes reach to the average of their state values. When v_i goes through the initialization update once with all its neighbors, it proceeds to the consensus phase.

A. Initialization Phase

During the initialization phase, the coupling weights are selected from the set of all real numbers. The update equations become

$$\begin{aligned} x_i^\alpha[k+1] &= x_i^\alpha[k] + a_{ij}[k](x_j^\alpha[k] - x_i^\alpha[k]) \\ &\quad + a_{i,\alpha\beta}[k](x_i^\beta[k] - x_i^\alpha[k]), \\ x_i^\beta[k+1] &= x_i^\beta[k] + a_{i,\alpha\beta}[k](x_i^\alpha[k] - x_i^\beta[k]), \end{aligned} \quad (3)$$

where $a_{i,\alpha\beta}[k] \in \mathbb{R}$ and $a_{ij}[k] \in \mathbb{R}$. Node v_i that will update its state variable is selected with equal probability $p_i = 1/N$. Node v_i selects a neighboring node v_j with probability $p_{j|i} = \frac{1}{|N_i \setminus S_i|}$ where S_i is the set of neighbors of v_i that it has gone through the initialization update. The selected nodes update their alpha and beta states using (3) at iteration k , while all the other nodes keep their states the same. The initialization phase of v_i ends when it has gone through the update (3) with all $v_j \in N_i$. If there are no neighbors left to go through the initialization, v_i selects a node from the set of neighbors that finished the initialization to start the consensus phase.

B. Consensus Phase

During the consensus phase, the nodes update their state variables to reach to the average of their initial values. Each node that has finished initialization merges their state variables into one by

$$x_i[k] = (x_i^\alpha[k] + x_i^\beta[k])/2, v_i \in N. \quad (4)$$

Node v_i is selected with equal probability $p_i = 1/N$. If v_i has gone through the initialization update with all its neighbors once, it selects a neighboring node v_j with equal probability $p_{j|i} = \frac{1}{|F_i|}$ where F_i is defined to be the set of neighbors of v_i that have finished initialization. Given that F_i is not empty, nodes v_i and v_j go through the consensus update defined in (1). If there is no neighbor that has finished initialization, the update is skipped. Although randomized gossip is selected, any averaging method can be used during the consensus phase. The initialization phase can be considered as a preprocessing step which transforms the initial values into another domain after which any distributed averaging method can be implemented in a privacy-preserving manner.

Theorem 1: Under Assumption 1, the proposed algorithm converges to the exact average of the initial values almost surely.

Proof. The coupling weights $a_{i,\alpha\beta}$ and $a_{ij} = a_{ji}$ are symmetric during the initialization for each node $v_i \in N$, $v_j \in N$ and $v_i \neq v_j$. Due to this, the sum of the network across the iterations is preserved. Using (3) and the symmetric weights,

$$\frac{1}{2N} \sum_{j=1}^N (x_j^\alpha[k] + x_j^\beta[k]) = \frac{1}{2N} \sum_{j=1}^N (x_j^\alpha[k+1] + x_j^\beta[k+1]).$$

After the initialization, the nodes merge their state variables into one through $x_i[k] = (x_i^\alpha[k] + x_i^\beta[k])/2$, which preserves the sum of the nodes' states. Using randomized gossip algorithm [19], the nodes will converge to the mean $\frac{1}{2N} \sum_{j=1}^N (x_j^\alpha[k] + x_j^\beta[k])$, which is equivalent to $\frac{1}{N} \sum_{j=1}^N x_j[0]$ due to the initial constraint $x_i^\alpha[0] + x_i^\beta[0] = 2x_i[0]$.

V. PRIVACY ANALYSIS

Following the convention in [12], privacy is defined as follows.

Definition 1: The privacy of the initial value $x_i[0]$ for any node v_i is preserved if an adversary cannot estimate the value of $x_i[0]$ with any guaranteed accuracy.

The privacy breach is explained by [13], which shows that if all the neighbors of v_i can be listened to by the passive adversary, the privacy cannot be established. The following assumption is made to prove privacy.

Assumption 2: One of the coupling weight $a_{ij}[k]$ is hidden from the adversary for each node $v_i \in N, v_j \in N, v_i \neq v_j$.

Theorem 2: Under Assumptions 1 and 2, the privacy as defined in Definition 1 will be achieved asymptotically as the variances of coupling weights go to infinity.

Proof. Let $x_i[0]$ be the initial value that the semi-honest adversary tries to estimate. The initial value can be found using the relation $2x_i[0] = x_i^\alpha[0] + x_i^\beta[0]$. Since $x_i^\alpha[0]$ is released and known by the adversary, estimating $x_i[0]$ is the same as estimating $x_i^\beta[0]$. There are two cases, which define the privacy of the initial value. The first one is the known coupling weight and the second one is the hidden coupling weight.

If the coupling weight can be captured by the adversary, $x_i^\alpha[k]$, $x_j^\alpha[k]$ and $a_{ij}[k]$ for $v_j \in N_i$ will be known and can be treated as a constant. The information leakage for this case can be defined as follows,

$$I(x_i^\alpha[k+1]; x_i^\beta[k] | x_i^\alpha[k], a_{ij}[k], x_j^\alpha[k]).$$

Using (3), the conditional mutual information becomes

$$I(a_{i,\alpha\beta}[k](x_i^\beta[k] - x_i^\alpha[k]); x_i^\beta[k] | x_i^\alpha[k]).$$

There should be no information leakage regarding $x_i^\beta[k]$ during this case since $x_i^\beta[k]$ is directly related to $x_i[0]$ due to the initial constraint. To establish the privacy, we will show that

$$\lim_{\sigma_{a_{i,\alpha\beta}[k]}^2 \rightarrow \infty} I(a_{i,\alpha\beta}[k](x_i^\beta[k] - x_i^\alpha[k]); x_i^\beta[k] | x_i^\alpha[k]) = 0. \quad (5)$$

The second case is when v_i contacts v_m with whom v_i shares a secret coupling weight a_{im} . Let $s[T]$ denote the sum $\sum_{k=1}^T a_{i,\alpha\beta}[k](x_i^\beta[k] - x_i^\alpha[k])$ that can be obtained by

$$s[k+1] = s[k] + x_i^\alpha[k+1] - a_{ij}[k](x_j^\alpha[k] - x_i^\alpha[k]), \quad (6)$$

where $s[0] = x_i^\alpha[0]$, T is the iteration in which v_i has finished going through the initialization update (3) with $v_j \in N_i$. If the adversary is able to listen all the communications of the targeted node, she can obtain the initial value using $s[T]$ and $x[T+1]$, since $x_i^\beta[T]$ will be disclosed with the first update of the consensus phase. Assumption 2 guarantees that $a_{i,\alpha\beta}[k](x_i^\beta[k] - x_i^\alpha[k])$ is blinded by $a_{ij}[k](x_j^\alpha[k] - x_i^\alpha[k])$ at least once. For this case the information leakage can be defined as follows,

$$I(x_i^\alpha[k+1]; a_{i,\alpha\beta}[k](x_i^\beta[k] - x_i^\alpha[k]) | x_i^\alpha[k]).$$

Using (3), the conditional mutual information becomes

$$I(a_{i,\alpha\beta}[k](x_i^\beta[k] - x_i^\alpha[k]) + a_{ij}[k](x_j^\alpha[k] - x_i^\alpha[k]); a_{i,\alpha\beta}[k](x_i^\beta[k] - x_i^\alpha[k]) | x_i^\alpha[k]).$$

For the privacy to be established, we will show that

$$\lim_{\substack{\sigma_{a_{ij}[k]}^2 \rightarrow \infty \\ \sigma_{a_{i,\alpha\beta}[k]}^2 \rightarrow \infty}} I(a_{i,\alpha\beta}[k](x_i^\beta[k] - x_i^\alpha[k]) + a_{ij}[k](x_j^\alpha[k] - x_i^\alpha[k]); a_{i,\alpha\beta}[k](x_i^\beta[k] - x_i^\alpha[k]) | x_i^\alpha[k]) = 0. \quad (7)$$

If (5) and (7) both hold, there will be no information leakage in the system about the initial values. The mutual information $I(x_i^\alpha[T]; x_i[0])$ will be zero since there will be no dependence between the alpha states and the initial value.

For simplicity of notation, let $W_{\alpha\beta} = a_{i,\alpha\beta}[k](x_i^\beta[k] - x_i^\alpha[k])$ and $W_{ij} = a_{ij}[k](x_j^\alpha[k] - x_i^\alpha[k])$. The iteration number $[k]$ is omitted in the equations and only written to explicitly state the next iteration or iteration 0.

First, it will be shown that for a fixed bounded variance $x_i^\alpha[0]$, the conditional mutual information $I(x_i^\alpha[k+1]; x_i^\beta[k] | x_i^\alpha[k])$ goes to zero as the variance of $a_{i,\alpha\beta}[k]$ goes to infinity. Let $x_i^\alpha[0]$ be a continuous random variable with $\sigma_{x_i^\alpha[0]}^2 < \infty$. Define $\gamma = \frac{1}{\sigma_{a_{i,\alpha\beta}[k]}^2}$ and $\bar{W}_{\alpha\beta} = \gamma a_{i,\alpha\beta}[k](x_i^\beta[k] - x_i^\alpha[k])$. The conditional mutual information can be written as follows,

$$I(x_i^\alpha[k+1]; x_i^\beta | x_i^\alpha, a_{ij}, x_j^\alpha) = I(W_{\alpha\beta}; x_i^\beta | x_i^\alpha).$$

The mutual information is invariant to scaling, that is,

$$I(\gamma W_{\alpha\beta}; \gamma x_i^\beta | \gamma x_i^\alpha) = I(\bar{W}_{\alpha\beta}; \gamma x_i^\beta | \gamma x_i^\alpha).$$

As the variance of $a_{i,\alpha\beta}[k]$ goes to infinity, the conditional mutual information will go to zero. Indeed, we have

$$\begin{aligned} \lim_{\sigma_{a_{i,\alpha\beta}}^2 \rightarrow \infty} I(\bar{W}_{\alpha\beta}; \gamma x_i^\beta | \gamma x_i^\alpha) &= \lim_{\gamma \rightarrow 0} I(\bar{W}_{\alpha\beta}; \gamma x_i^\beta | \gamma x_i^\alpha) \\ &= I(\bar{W}_{\alpha\beta}; 0) = 0. \end{aligned}$$

The second case is the update without the knowledge of a_{ij} . Define $\beta = \frac{1}{\sigma_{a_{ij}[k]}^2}$ and $\bar{W}_{ij} = \beta a_{ij}[k](x_i^\beta[k] - x_i^\alpha[k])$. Mutual information is invariant to scaling, so that,

$$\begin{aligned} I(x_i^\alpha[k+1]; W_{\alpha\beta}|x_i^\alpha) &= I(\gamma\beta x_i^\alpha[k+1]; \gamma\beta W_{\alpha\beta}|\gamma\beta x_i^\alpha) \\ &= I(\beta\bar{W}_{\alpha\beta} + \gamma\bar{W}_{ij}; \beta\bar{W}_{\alpha\beta}|\gamma\beta x_i^\alpha). \end{aligned}$$

When the variances of both coupling weights go to infinity, the conditional mutual information will go to zero.

$$\begin{aligned} \lim_{\substack{\sigma_{a_{ij}}^2 \rightarrow \infty \\ \sigma_{a_{i,\alpha\beta}}^2 \rightarrow \infty}} I(\beta\bar{W}_{\alpha\beta} + \gamma\bar{W}_{ij}; \beta\bar{W}_{\alpha\beta}|\gamma\beta x_i^\alpha) &= \\ \lim_{\substack{\gamma \rightarrow 0 \\ \beta \rightarrow 0}} I(\beta\bar{W}_{\alpha\beta} + \gamma\bar{W}_{ij}; \beta\bar{W}_{\alpha\beta}|\gamma\beta x_i^\alpha) &= I(0; 0) = 0. \end{aligned}$$

Let T be the iteration at which the initialization has ended. $\mathbf{x}^\alpha[T]$ represents the vector of alpha values obtained starting from $\mathbf{x}[0]$. During the consensus phase, let W^k denote the information obtained at each iteration to deduce $\mathbf{x}[0]$ with $k = \{1, 2, \dots, K\}$ where K is the total number of iterations. The final mutual information can be represented as $I(\mathbf{x}[0]; W^k)$. Fixing the coupling weights, enables to find a function $F^k(\mathbf{x}^\alpha[T]) = W^k$ that will take the $\mathbf{x}^\alpha[T]$ as input and will create the output W^k . The random variables will create a Markov chain $\mathbf{x}[0] \rightarrow \mathbf{x}^\alpha[T] \rightarrow W^k$ for $k = \{1, 2, \dots, K\}$. The data processing inequality [24] shows that

$$I(\mathbf{x}[0]; W^k) \leq I(\mathbf{x}[0]; \mathbf{x}^\alpha[T]) = 0, \quad k = 1, \dots, K.$$

No clever manipulation of data can increase the mutual information. Thus, given that there is at least one coupling weight a_{ij} that is hidden from the adversary, the semi-honest adversary cannot estimate the initial value of node v_i with any guaranteed accuracy.

A similar privacy proof can be done for the eavesdropper case. If Assumption 2 holds, the eavesdropper will not have any range to estimate the initial value with any accuracy as the variances of coupling weights go to infinity.

In practice, it is not possible to have infinitely large variances. However, given that $\frac{\sigma_{a_{i,\alpha\beta}[0]}^2}{\sigma_x^2[0]} = \frac{\sigma_{a_{ij}[0]}^2}{\sigma_x^2[0]} = \frac{\sigma_{x^\alpha[0]}^2}{\sigma_x^2[0]} = 100$ (the range of the coupling weights and alpha states are approximately 10 times the range of the input $x[0]$), the information leakage $I(x_i^\alpha[1]; x_i[0]|x_i^\alpha[0])$ is only 0.01 bits, which can be considered small.

VI. EXPERIMENTS

To demonstrate the performance of the proposed approach, a 5-node circular graph is generated. In practical applications, the transmission power of the nodes in the graph can be arranged such that the Assumptions 1 and 2 are satisfied; the graph is connected and one of the coupling weight a_{ij} is hidden from a semi-honest adversary in the network. To establish the privacy against an eavesdropper who can tap arbitrarily any channel, homomorphic encryption [25] or secret sharing schemes with channel encryption [20] can be used.




Fig. 1. Conditional mutual information $I(x^\alpha[1]; x[0]|x^\alpha[0])$ plot for different $x^\alpha[0]$ and $a_{i,\alpha\beta}[0]$ variances sampled from uniform distributions.

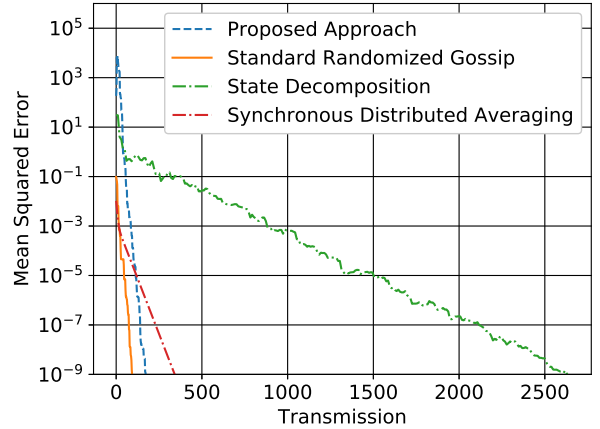


Fig. 2. Convergence rate plot for the proposed approach, state decomposition [12], synchronous distributed averaging with maximum-degree weights [19] and the standard randomized gossip.

Two experiments have been done to assess the provided privacy and convergence properties. The provided privacy represented by $I(x^\alpha[1]; x[0]|x^\alpha[0])$ is plotted against increasing $x^\alpha[0]$ and $a_{i,\alpha\beta}[0]$ variances in Fig. 1. In words, the term $I(x^\alpha[1]; x[0]|x^\alpha[0])$ represents the first case in privacy analysis which assesses how much information about the initial value is leaked with the release of the first alpha state given the knowledge of the initial alpha states and the coupling weights. A lower result means there is less dependency between $x^\alpha[1]$ and $x[0]$, which suggests that less information about the initial values can be gained by the knowledge of $x^\alpha[1]$.

The conditional mutual information [26] is estimated using the non-parametric mutual information toolbox (NPEET) [27] which uses k-nearest neighbour entropy estimates [28]. The variances $\sigma_{a_{i,\alpha\beta}[0]}^2$ and $\sigma_{x^\alpha[0]}^2$ are increased with respect to the unit variance $\sigma_{x[0]}^2$. Both $a_{i,\alpha\beta}[0]$ and $x^\alpha[0]$ are sampled from uniform distributions with variances (1, 4, 9, 16, 25, 36, 49, 64, 81, 100). For each combination, 10^4 independent samples are drawn and the conditional mutual information is estimated. The experiment is repeated over all variances and the averaged results are plotted in Fig. 1. When the variances increase, the conditional mutual information approaches zero. This shows that the new alpha states contain less information about the initial values with the increasing variance. A similar result is observed for the second case in the privacy analysis where the coupling weight a_{ij} is unknown.

The convergence rate is assessed with the second experiment. The substate $x^\alpha[0]$ and the coupling weights are selected from the interval $[-5, 5]$, while the initial values are selected from the interval $[1, 2]$ uniformly at random. As shown in Fig. 2, the convergence rate of the state decomposition approach is lower than the synchronous distributed averaging with maximum-degree weight [19] without the privacy-preserving attribute, because of the increase in the number of nodes due to the state decomposition. In the proposed approach, the nodes merge their state variables into one after the initialization. The convergence rate, in this case, is the same as the randomized gossip algorithm. If the update equations in (2) are used during the consensus phase, the convergence rate is lowered significantly. The privacy guarantee can be given in the initialization phase. Hence, reduction in the convergence rate can be avoided.

VII. CONCLUSIONS

In this paper, an asynchronous privacy-preserving average consensus algorithm is proposed using the state decomposition approach. An information-theoretic privacy analysis is done, which promises to preserve the privacy of initial values against an eavesdropper and a semi-honest adversary given that there is at least one coupling weight hidden from the adversary. The proposed approach converges to the exact average of the initial values, while keeping the convergence rate the same as the standard randomized gossip algorithm. The privacy can be guaranteed through the addition of the noise for a bounded amount of time and without sacrificing the convergence rate. Moreover, the algorithm does not need a trusted third party. Comparison of privacy-preserving schemes with respect to the information leakage is left for future work.

REFERENCES

- [1] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 215–233, Jan 2007.
- [2] R. Olfati-Saber, "Distributed kalman filter with embedded consensus filters," in *Proceedings of the 44th IEEE Conference on Decision and Control*, Dec 2005, pp. 8179–8184.
- [3] R. Olfati-Saber, "Flocking for multi-agent dynamic systems: algorithms and theory," *IEEE Transactions on Automatic Control*, vol. 51, no. 3, pp. 401–420, March 2006.
- [4] Y. Cao and W. Ren, "Distributed coordinated tracking with reduced interaction via a variable structure approach," *IEEE Transactions on Automatic Control*, vol. 57, no. 1, pp. 33–48, Jan 2012.
- [5] A. Jadbabaie, Jie Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Transactions on Automatic Control*, vol. 48, no. 6, pp. 988–1001, June 2003.
- [6] J. Lin, A. S. Morse, and B. D. O. Anderson, "The multi-agent rendezvous problem," in *42nd IEEE International Conference on Decision and Control (IEEE Cat. No.03CH37475)*, Dec 2003, vol. 2, pp. 1508–1513 Vol.2.
- [7] C. Zhao, J. Chen, J. He, and P. Cheng, "Privacy-preserving consensus-based energy management in smart grids," *IEEE Transactions on Signal Processing*, vol. 66, no. 23, pp. 6162–6176, Dec 2018.
- [8] M. Kishida, "Encrypted average consensus with quantized control law," in *2018 IEEE Conference on Decision and Control (CDC)*, Dec 2018, pp. 5850–5856.
- [9] R. Lazeretti, S. Horn, P. Braca, and P. Willett, "Secure multi-party consensus gossip algorithms," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2014, pp. 7406–7410.
- [10] 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.
- [11] J. He, L. Cai, C. Zhao, and P. Cheng, "Privacy-preserving average consensus: Privacy analysis and optimal algorithm design," *IEEE Transactions on Signal and Information Processing over Networks*, vol. PP, 09 2016.
- [12] Y. Wang, "Privacy-preserving average consensus via state decomposition," *IEEE Transactions on Automatic Control*, 2019.
- [13] Y. Mo and R. M. Murray, "Privacy preserving average consensus," *IEEE Trans. Automat. Contr.*, vol. 62, no. 2, pp. 753–765, 2017.
- [14] P. Paillier, "Public-key cryptosystems based on composite degree residuosity classes," in *Proceedings of the 17th International Conference on Theory and Application of Cryptographic Techniques*, Berlin, Heidelberg, 1999, EUROCRYPT'99, pp. 223–238, Springer-Verlag.
- [15] Z. Huang, R. Hu, E. Chan-Tin, and Y. Gong, "DP-ADMM: admm-based distributed learning with differential privacy," *CoRR*, vol. abs/1808.10101, 2018.
- [16] C. Dwork, "Differential privacy," in *Automata, Languages and Programming*, M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, Eds., Berlin, Heidelberg, 2006, pp. 1–12, Springer Berlin Heidelberg.
- [17] C. Dwork, M. Naor, T. Pitassi, and G. Rothblum, "Differential privacy under continual observation," Jun 2010, pp. 715–724.
- [18] Q. Li, I. Cascudo, and M. G. Christensen, "Privacy-preserving distributed average consensus based on additive secret sharing," in *2019 27th European Signal Processing Conference (EUSIPCO)*, Sep. 2019, pp. 1–5.
- [19] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized gossip algorithms," *IEEE/ACM Trans. Netw.*, vol. 14, no. SI, pp. 2508–2530, June 2006.
- [20] Q. Li, I. Cascudo, and M. G. Christensen, "Privacy-preserving distributed average consensus based on additive secret sharing," in *2019 27th European Signal Processing Conference (EUSIPCO)*, IEEE, 2019, pp. 1–5.
- [21] F. Hanzely, J. Konečný, N. Loizou, P. Richtárik, and D. Grishchenko, "Privacy preserving randomized gossip algorithms," *arXiv preprint arXiv:1706.07636*, 2017.
- [22] A. Shamir, "How to share a secret," *Commun. ACM*, vol. 22, no. 11, pp. 612–613, Nov. 1979.
- [23] F. Fagnani and S. Zampieri, "Asymmetric randomized gossip algorithms for consensus," *IFAC Proceedings Volumes*, vol. 41, no. 2, pp. 9051 – 9056, 2008, 17th IFAC World Congress.
- [24] T. M. Cover and J. A. Thomas, *Elements of information theory*, John Wiley & Sons, 2012.
- [25] M. Ruan, M. Ahmad, and Y. Wang, "Secure and privacy-preserving average consensus," in *Proceedings of the 2017 Workshop on Cyber-Physical Systems Security and Privacy*, 2017, pp. 123–129.
- [26] M. Vejmelka and M. Paluš, "Inferring the directionality of coupling with conditional mutual information," *Phys. Rev. E*, vol. 77, pp. 026214, Feb 2008.
- [27] G. V Steeg, "Non-parametric entropy estimation toolbox (npeet)," 2000.
- [28] A. Kraskov, H. Stögbauer, and P. Grassberger, "Estimating mutual information," *Physical Review E*, vol. 69, no. 6, Jun 2004.