Privacy-Preserving Distributed Maximum Consensus

Naveen K. D. Venkategowda, *Member, IEEE,* and Stefan Werner, *Senior Member, IEEE*

***Abstract*—We propose a privacy-preserving distributed maximum consensus algorithm where the local state of the agents and identity of the maximum state owner is kept private from adversaries. To that end, we reformulate the maximum consensus problem over a distributed network as a linear program. This optimization problem is solved in a distributed manner using the alternating direction method of multipliers (ADMM) and perturbing the primal update step with Gaussian noise. We define the privacy of an agent as the estimation error of its local state at the adversary and obtain theoretical bounds on the privacy loss for the proposed method. Further, we prove that the proposed algorithm converges to the maximum value at all agents. In addition to the analytical results, we illustrate the convergence speed and privacy-accuracy trade-off through numerical simulations.**

I. Introduction

Consensus in a multi-agent system such as average consensus and maximum/minimum consensus is required in distributed computing, optimization, control and robotics [1]–[3]. Consensus algorithms are based on local computation and exchanging information with neighboring agents to reach a network-wide agreement on the desired value. These algorithms require the agents to share their local state with the neighboring agents, which may result in loss of privacy. For example, in smart grids where multiple generators must reach a consensus on the cost while not revealing their information about individual generation [4]. In the multi-agent rendezvous problem, a group of agents agree to rendezvous at a particular location but may not want to disclose their initial locations [5]. In many instances the identity of agents that own the consensus value needs to be private. For instance, the identity of the leader in distributed control with leader-follower multi-agent network [6] and the identity of the cluster head in sensor networks [7] must be private to safeguard them from attacks.

Works in [8]–[10] propose privacy-preserving average consensus algorithms where agents add noise to their state updates and message-generating functions. In contrast to [8], which considers differential privacy, [9] considers (α, β)-data-privacy that captures the maximum disclosure probability and estimation accuracy, whereas [10] assumes the privacy metric to be the covariance of the maximum likelihood estimate of the local states at the eavesdropper. Secure multiparty computation based methods have been proposed in [11]–[13] for average consensus with privacy guarantees. In [14] a deterministic approach exploiting homomorphic cryptography to enforce secrecy during interaction between nodes was proposed for distributed consensus algorithms. It was shown in [15] that for a network with certain topological restrictions, average consensus can be achieved and the agents' local data is completely unobservable from the data received at a given agent.

Although distributed maximum consensus has been investigated under various settings [16]–[18], there is a lack of studies on maximum consensus with privacy guarantees. Authors in [19] proposed a differentially private maximum consensus algorithm based on adding Laplacian noise to the initial states and proved that exact consensus and differential privacy cannot be guaranteed simultaneously. In [20] a privacy-preserving maximum consensus algorithm was proposed where agents broadcast random data before transmitting their actual states. Though [20] characterizes the probability of maximum state owner's identity being revealed, it does not quantify the privacy leakage of other agents.

In this paper, we characterize the privacy metric as the covariance of a minimum mean square error (MMSE) estimate of an agent's state at the adversary. We propose a privacy-preserving distributed maximum consensus algorithm in which the adversary will not be able to infer exact local state of the agents. We first transform the maximum consensus problem into a simple linear program that is solved in a distributed manner using the alternating direction method of multipliers (ADMM). To endow privacy, every agent uses a random initialization which is unknown to other agents and the primal update step is perturbed with zero-mean Gaussian noise whose variance decays at each message-sharing step. We obtain theoretical bounds on the privacy leakage at the agents and prove that the proposed algorithm converges. Further, we illustrate the convergence and privacy-accuracy trade-off through numerical simulations.

II. Problem Formulation

We consider a connected network of $L \in \mathbb{N}$ agents modeled as an undirected graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ where the set of vertices $\mathcal{V} = \{1, \ldots, L\}$ corresponds to the agents and the edge set $\mathcal{E}$ represents the communication links between the pairs of agents. Agent $i \in \mathcal{V}$ can communicate with its neighbors whose indices are in the set $\mathcal{N}_i$ with cardinality $|\mathcal{N}_i|$. By convention, $\mathcal{N}_i$ does not include the agent i itself.

Let a_i denote the local state at agent i and a^* indicate the maximum value, i.e. $a^* = \max_i a_i$. The maximum consensus algorithm ensures that every agent attains consensus on the maximum value among the agents by local computation and exchanging information with their neighbors. The distributed maximum consensus algorithm is given by

$$x_i(k+1) = \max_{j \in \mathcal{N}_i} (x_i(k), x_j(k)), \quad i \in \mathcal{V}, \tag{1}$$

with initial value chosen as $x_i(0) = a_i$. Further, there exists a finite T such that $x_i(k) = a^*$ for all $k \geq T$ and $i = 1, \ldots, L$.

This work was partly supported by the Research Council of Noway and Academy of Finland under Grant 296849. The authors are with the Department of Electronic Systems, Norwegian University of Science and Technology, Trondheim, Norway. Email: {naveen.dv, stefan.werner}@ntnu.no

It can be observed that every agent must share its local state with the neighboring agents. However, in many applications local states $\{a_i\}_{i=1}^L$ must remain private. An adversary can infer the local state of other agents using the data received from its neighbors. Further, the identity of the agent with maximum value a^* will be revealed, which may be undesirable. Therefore, the objective of this paper is to develop a distributed maximum consensus algorithm that attains maximum consensus and preserves the privacy of local state.

III. Privacy-Preserving Distributed Maximum Consensus

The maximum consensus algorithm (1) has a nonlinear update step in contrast to the average consensus algorithms. Hence, we cannot directly employ the perturbation-based privacy-preserving approaches to maximum consensus as the local states at the agents may diverge. However, we overcome this problem by first reformulating the maximum consensus as the following optimization problem:

$$\begin{aligned} a^* = \arg\min_{x} \quad & x \\ \text{s. t.} \quad & x \geq a_i, \quad i = 1, \ldots, L. \end{aligned} \tag{2}$$

Let x_i denote the local value of x at agent i. In order to solve (2) in a distributed manner, we next recast it as a linear program given by

$$\begin{aligned} \min_{\{x_i\}} \quad & \frac{1}{L}\sum_{i=1}^{L} x_i \\ \text{s.t.} \quad & x_i \geq a_i, \quad i = 1, \ldots, L \\ & x_i = x_j, \quad j \in \mathcal{N}_j, \quad \forall i \in \mathcal{V}. \end{aligned} \tag{3}$$

The equality constraints enforce local consensus across each agent's neighborhood. Now solving the maximum consensus problem is equivalent to addressing the optimization problem (3) since both (2) and (3) have an identical solution, i.e. the solution of (3) $x_i^* = a^*$, for all i. Since the local objective functions are affine and the agents' local constraints constitute the sensitive information, we cannot employ existing distributed algorithms [21]–[24] to guarantee privacy.

By introducing the auxiliary local variables $\mathcal{Z} = \{\{z_i^j\}_{j\in\mathcal{N}_i}\}_{i=1}^L$ and using the indicator function, defined as $\mathcal{I}_a(y) = 0$, if $y \geq a$ and ∞ otherwise, to enforce the inequality constraint, we rewrite the problem in (3) as

$$\begin{aligned} \min_{\{x_i, y_i, z_i^j\}} \quad & \frac{1}{L}\sum_{i=1}^{L} x_i + \frac{1}{L}\sum_{i=1}^{L} \mathcal{I}_{a_i}(y_i) \\ \text{s.t.} \quad & x_i = y_i, \quad i = 1, \ldots, L \\ & x_i = z_i^j, \; x_j = z_i^j, \quad i \neq j, j \in \mathcal{N}_i, \forall i \in \mathcal{V}. \end{aligned} \tag{4}$$

The auxiliary variables $\mathcal{Z}$ are used to obtain an equivalent representation of the constraints in (3) and will be eliminated eventually. As the objective function in (4) is separable, the ith agent can independently compute the optimal x_i^* by relying on the ADMM technique to solve (4) in a distributed manner [25]. For this purpose, the augmented Lagrangian for (4) with quadratic penalty for constraint violations is expressed as

$$\begin{aligned} &\mathcal{L}_\rho(\{x_i\}_{i=1}^L, \{y_i\}_{i=1}^L, \mathcal{Z}, \mathcal{M}) = \\ &\sum_{i=1}^{L}\Big(\frac{x_i}{L} + \frac{1}{L}\mathcal{I}_{a_i}(y_i) + u_i(x_i - y_i) + \frac{\rho_y}{2}(x_i - y_i)^2\Big) \\ &+ \sum_{i=1}^{L}\sum_{j\in\mathcal{N}_i}\Big(\mu_i^j(x_i - z_i^j) + \lambda_i^j(x_j - z_i^j)\Big) \\ &+ \frac{\rho_z}{2}\sum_{i=1}^{L}\sum_{j\in\mathcal{N}_i}\Big((x_i - z_i^j)^2 + (x_j - z_i^j)^2\Big), \end{aligned} \tag{5}$$

where $\mathcal{M} := \{u_i, \{\mu_i^j\}_{j\in\mathcal{N}_i}, \{\lambda_i^j\}_{j\in\mathcal{N}_i}\}_{i=1}^L$ are the Lagrange multipliers, and ρ_y and ρ_z are the penalty parameters associated with first and second constraints in (4), respectively.

To solve the minimization problem (4) in a distributed fashion, $\mathcal{L}_\rho$ is minimized with respect to the primal variables $\{x_i\}_{i=1}^L$, $\{y_i\}_{i=1}^L$, and auxiliary variables $\mathcal{Z}$ alternately with the other two sets of variables fixed. Then, the Lagrange multipliers in $\mathcal{M}$ are updated via dual gradient-ascent iterations [25]. It can be seen that the Lagrangian in (5) is separable in x_i and y_i, and the penalty terms can be simplified as $u_i(x_i - y_i) + \frac{\rho_y}{2}(x_i - y_i)^2 = \frac{\rho_y}{2}(x_i - y_i + \tilde{u}_i)^2 - \frac{\rho_y}{2}\tilde{u}_i^2$, where $\tilde{u}_i = u_i/\rho_y$. By using the Karush-Kuhn-Tucker optimality conditions for (4) and setting $v_i(k) = 2\sum_{j\in\mathcal{N}_i}\lambda_i^j(k)$, it can be shown that the Lagrange multipliers $\{\mu_i^j\}_{j\in\mathcal{N}_i}$ and the auxiliary variables $\mathcal{Z}$ are eliminated [26]. Therefore, the distributed ADMM algorithm to solve (4) reduces to the following iterative updates at the ith agent

$$\begin{aligned} x_i(k+1) = \arg\min_{x_i} \Bigg\{ &\frac{1}{L}x_i + \frac{\rho_y}{2}(x_i - y_i(k) + \tilde{u}_i(k))^2 \\ &+ v_i(k)x_i + \rho_z \sum_{j\in\mathcal{N}_i}\left(x_i - \frac{x_i(k) + x_j(k)}{2}\right)^2 \Bigg\} \end{aligned} \tag{6}$$

$$y_i(k+1) = \arg\min_{y_i} \frac{\mathcal{I}_{a_i}(y_i)}{L} + \frac{\rho_y}{2}(x_i(k+1) - y_i + \tilde{u}_i(k))^2 \tag{7}$$

$$\tilde{u}_i(k+1) = \tilde{u}_i(k) + x_i(k+1) - y_i(k+1) \tag{8}$$

$$v_i(k+1) = v_i(k) + \rho_z \sum_{j\in\mathcal{N}_i}[x_i(k+1) - x_j(k+1)] \tag{9}$$

where k is the iteration index and $\{y_i(0)\}_{i=1}^L$, $\{\tilde{u}_i(0)\}_{i=1}^L$, $\{v_i(0)\}_{i=1}^L$ are set to zero.

Next, we derive closed-form solutions to address the subproblems in (6) and (7). We can see that (6) is an unconstrained quadratic optimization problem. Thus, by computing the gradient of the objective function in (6) and equating it to zero, the optimal update at the ith node is obtained as

$$\begin{aligned} x_i(k+1) = &\frac{\rho_z}{\rho_y + 2|\mathcal{N}_i|\rho_z}\sum_{j\in\mathcal{N}_i}(x_i(k) + x_j(k)) \\ &+ \frac{\rho_y}{\rho_y + 2|\mathcal{N}_i|\rho_z}[y_i(k) - u_i(k)] - \frac{v_i(k) + (1/L)}{\rho_y + 2|\mathcal{N}_i|\rho_z}. \end{aligned} \tag{10}$$

The second minimization step (7) in the algorithm, $\tilde{f}(x_i(k+1) + \tilde{u}_i(k)) = \arg\min_{y_i}\mathcal{I}_{a_i}(y_i) + \frac{\rho_y L}{2}(x_i(k+1) - y_i + \tilde{u}_i(k))^2$, is the proximal operator of the indicator function of a closed nonempty convex set $\{y_i \in \mathbb{R} \,|\, y_i \geq a_i\}$. It is known that the proximal operator of the indicator function $\mathcal{I}_{a_i}(y_i)$ is the

projection of $x_i(k+1)+\tilde{u}_i(k)$ onto set $\{y_i \in \mathbb{R} \,|\, y_i \geq a_i\}$. Hence, the update step for y_i is given by

$$y_i(k+1) = \max(x_i(k+1)+\tilde{u}_i(k), a_i). \tag{11}$$

It is apparent from (9) and (10) that the agents need to transmit $x_i(k)$ to their neighboring nodes to compute the network-wide maximum of the local state. However, this sharing process aids curious agents to infer the local data that the agents want to keep private. To prevent the adversary from knowing the private state, the messages shared by every agent are perturbed before transmission. The ith agent chooses a random initial point $x_i(0)$ with zero-mean Gaussian distribution and variance σ_x^2, i.e. $x_i(0) \sim \mathcal{N}(0, \sigma_x^2)$. This random number is kept secret and not shared with other agents.

Next, at the kth ADMM iteration, the ith agent generates a random variable $n_i(k)$ with normal distribution $\mathcal{N}(0, \sigma_n^2(k))$ and $\mathbb{E}[n_i(k)n_j(l)] = 0$ for $k \neq l$ and $i \neq j$, where $\sigma_n^2(k)$ is the variance of the perturbation noise $n_i(k)$ that decreases with iteration index k, i.e., $\sigma_n^2(k+1) < \sigma_n^2(k)$. The perturbed message transmitted from ith agent to all its neighboring agents $j \in \mathcal{N}_i$ is expressed as $\tilde{x}_i(k) = x_i(k) + n_i(k)$. Therefore, the dual variable updated through message sharing is given by

$$v_i(k+1) = v_i(k) + \rho_z \sum_{j \in \mathcal{N}_i} [\tilde{x}_i(k+1) - \tilde{x}_j(k+1)]. \tag{12}$$

Collating the steps in (10)–(12), the proposed privacy-preserving distributed maximum consensus is summarized in Algorithm 1. In the next section, we show that the proposed algorithm with secret random initialization and message perturbation is privacy-preserving.

Algorithm 1 Privacy-Preserving Maximum Consensus

At all agents $i \in \mathcal{V}$, initialize $x_i(0) \sim \mathcal{N}(0, \sigma_x^2)$, $y_i(0) = 0, \tilde{u}_i(0) = 0, v_i(0) = 0$
for $i = 1, 2, \ldots, L$ **do**
  for $k = 0, 1, \ldots$ **do**
    if $k = 0$ **then**
      $x_i(1) = \frac{\rho_z}{\rho_y + 2|\mathcal{N}_i|\rho_z} x_i(0) + \frac{1}{L(\rho_y + 2|\mathcal{N}_i|\rho_z)}$
    else
      Update $x_i(k+1)$ as in (10)
    end if
    Update $y_i(k+1) = \max(x_i(k+1) + \tilde{u}_i(k), a_i)$
    Update $\tilde{u}_i(k+1) = \tilde{u}_i(k) + x_i(k+1) - y_i(k+1)$
    Generate $n_i(k+1) \sim \mathcal{N}(0, \sigma_n^2(k+1))$
    Share $\tilde{x}_i(k+1) = x_i(k+1) + n_i(k+1)$ with agents in neighborhood $\mathcal{N}_i$
    Update dual variable $v_i(k+1) = v_i(k) + \rho_z \sum_{j \in \mathcal{N}_i} [\tilde{x}_i(k+1) - \tilde{x}_j(k+1)]$
  end for
end for

A. Privacy Guarantees

The information at the adversary at time k to estimate the state of agent i is defined as

$$\mathcal{X}_i(k) = \{\tilde{x}_i(1), \tilde{x}_i(2), \ldots, \tilde{x}_i(k)\}. \tag{13}$$

We assume that the adversary has the knowledge of the network, ADMM penalty, and perturbation noise variance. Hence, Algorithm 1 is valid for any honest-but-curious adversary, which can be an external eavesdropper or an agent in the network. The adversary computes an MMSE estimate of the local state a_i given the information $\mathcal{X}_i(k)$. Let us denote $\hat{a}_i(k)$ as the MMSE estimate and $P_i(k)$ as the estimator error covariance. Similar to [10] and [27], the privacy measure of node i is $P_i(k)$ and the privacy of node i is breached if $P_i(k) = 0$.

Theorem 1. *Algorithm 1 is privacy-preserving, i.e. $P_i(k) > 0$ with $P_i(k)$ bounded as $P_i(k) \geq Q_i(k)$ for $k < \infty$, $i = 1, \ldots, L$, and $Q_i(k)$ is given by*

$$Q_i(k) = \frac{Q_i(k-1)(\sigma_i^2(k) + (1-\gamma_i(k))\sigma_i^2(k))}{Q_i(k-1) + \sigma_i^2(k) + (1-\gamma_i(k))\sigma_i^2(k)}, \tag{14}$$

where $\gamma_i(k) = 1$ if $x_i(k) + \tilde{u}_i(k-1) \leq a_i$ and zero otherwise, and $\sigma_i^2(k) = \left(\frac{\rho_y + 2\rho_z|\mathcal{N}_i|}{2\rho_y}\right)^2 \sigma_n^2(k)$.

Proof. Since $\tilde{u}_i(0) = v_i(0) = 0$ for all i, we express (8) and (9) as $v_i(k) = \rho_z \sum_{l=1}^{k} \sum_{j \in \mathcal{N}_i} (x_i(l) - x_j(l))$ and $\tilde{u}_i(k) = \sum_{l=1}^{k} (x_i(l) - y_i(l))$. Substituting the above expressions and $y_i(k) = \max(x_i(k) + \tilde{u}_i(k-1), a_i)$ in (10), we obtain the observation dynamics at the adversary as

$$\tilde{x}_i(k+1) = -\alpha_y \tilde{x}_i(k) + 2\alpha_y \max(x_i(k) + \tilde{u}_i(k-1), a_i) + w_i(k) + n_i(k+1), \tag{15}$$

where $\alpha_y = \frac{\rho_y}{\rho_y + 2|\mathcal{N}_i|\rho_z}$, $\alpha_z = \frac{\rho_z}{\rho_y + 2|\mathcal{N}_i|\rho_z}$, and $w_i(k) = 2\alpha_z \sum_{j \in \mathcal{N}_i} \tilde{x}_j(k) + \sum_{l=1}^{k-1} \big(\alpha_y y_i(l) - \alpha_z \sum_{j \in \mathcal{N}_i} (\tilde{x}_i(l) - \tilde{x}_j(l))\big)$. Since the adversary has access to $\tilde{x}_i(l)$ for $l = 1, \ldots, k$, the new data available at time k to estimate a_i is given by

$$\tilde{x}_i(k+1) = \begin{cases} x_i(k) + \tilde{n}_i(k+1), & \gamma_i(k) = 0 \\ a_i + \tilde{n}_i(k+1), & \gamma_i(k) = 1 \end{cases} \tag{16}$$

where $\gamma_i(k) = 1$ if $x_i(k) + \tilde{u}_i(k-1) \leq 0$ and zero otherwise, and measurement noise $\tilde{n}_i(k+1) = \frac{\rho_y + 2\rho_z|\mathcal{N}_i|}{2\rho_y}(w_i(k) + n_i(k+1))$ with variance $R_i(k) = \mathbb{E}[\tilde{n}_i^2(k)]$. Hence the observation process at the adversary can be viewed as an stochastic event-triggered estimation of parameter a_i. This is equivalent to the model considered in [28] and [29] for remote estimation with open loop scheduling. From [28] and [29], an MMSE estimator for a_i for the model in (16) is given by

$$\hat{a}_i(k) = (1 - K_i^f(k))\hat{a}_i(k-1) + \gamma_i(k) K_i^f(k) \tilde{x}_i(k), \tag{17}$$

where filtering gain $K_i^f(k) = P_i(k-1)(P_i(k-1) + R_i(k) + (1-\gamma_i(k))R_i(k))^{-1}$ and $P_i(k) = \mathbb{E}[(\hat{a}_i(k) - a_i)^2 | \mathcal{X}_i(k)]$ is the estimation error covariance that follows $P_i(k) = (1 - K_i^f(k))P_i(k-1)$. Assuming that the adversary has knowledge of $\gamma_i(k)$ and $\sigma_i^2(k) = \left(\frac{\rho_y + 2\rho_z|\mathcal{N}_i|}{2\rho_y}\right)^2 \sigma_n^2(k) < R_i(k) = \mathbb{E}[\tilde{n}_i^2(k)]$, from [28] and [29] we can lower bound the estimation error at the adversary $P_i(k)$ with $Q_i(k) > 0$ that follows

$$Q_i(k) = \frac{Q_i(k-1)(\sigma_i^2(k) + (1-\gamma_i(k))\sigma_i^2(k))}{Q_i(k-1) + \sigma_i^2(k) + (1-\gamma_i(k))\sigma_i^2(k)}$$

with $Q_i(0) = \sigma_x^2$ and $\sigma_i^2(k) = \left(\frac{\rho_y + 2\rho_z|\mathcal{N}_i|}{2\rho_y}\right)^2 \sigma_n^2(k)$, $\forall i$. □

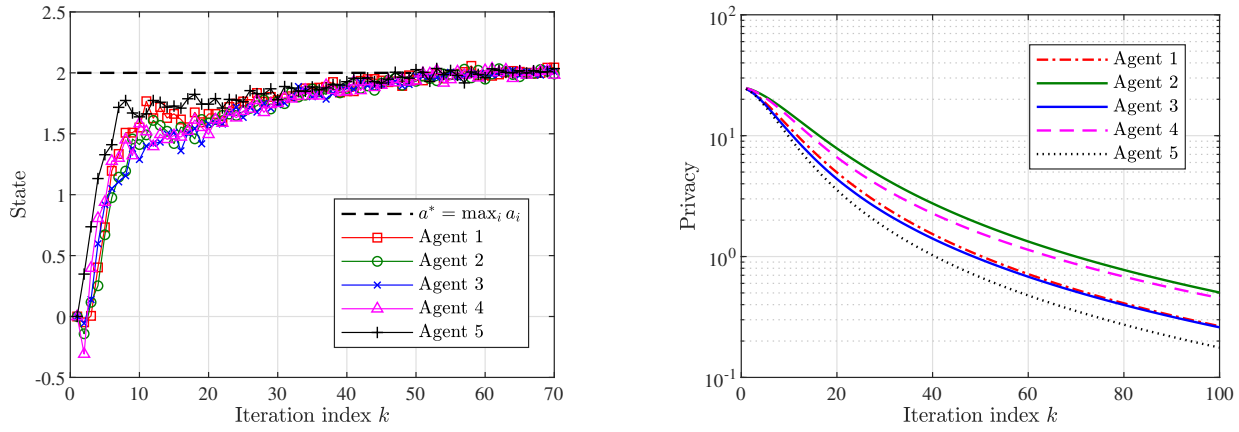

(a) (b)

Fig. 1. (a) Convergence to maximum value $a^* = 2$ for $\sigma_x^2 = \sigma_n^2(1) = 10$, (b) Privacy vs iteration index for perturbation noise variance $\sigma_n^2(1) = 10$.

B. Convergence

The primal variable $x_i(k)$ is perturbed with random variable $n_i(k)$ with normal distribution $\mathcal{N}(0, \sigma_n^2(k))$ and $\mathbb{E}[n_i(k)n_j(l)] = 0$ for $k \neq l$ and $i \neq j$, with the variance $\sigma_n^2(k)$ of the perturbation noise $n_i(k)$ decreasing with k. Further, the objective function in (4) is convex. Hence, from [30, Theorem 5], for any fixed number of iterations K, we have

$$\mathbb{E}\Big[\frac{\mathbf{1}_{2L}^{\mathsf{T}}\bar{\mathbf{x}}(K)}{L} - a^*\Big] \leq \frac{\rho_z \left\|\bar{\mathbf{x}}(0)\right\|_{\mathbf{L}_-}^2}{2K} + \frac{\rho_z \left\|\bar{\mathbf{x}}(0) - a^*\mathbf{1}_{2L}\right\|_{\mathbf{L}_+}^2}{2K} + \frac{\rho_z L \sigma_n^2(1)\phi_{\max}^2(\mathbf{L}_+)}{K(1-D)\phi_{\min}(\mathbf{L}_-)}, \tag{18}$$

where $\mathbf{1}_{2L} \in \mathbb{R}^{2L}$ denotes the vector of ones, $\bar{\mathbf{x}}(K) = \frac{1}{K}\sum_{k=1}^{K}\hat{\mathbf{x}}(k)$, $\hat{\mathbf{x}}(k) = [x_1(k), y_1(k), \ldots, x_L(k), y_L(k)]^{\mathsf{T}}$, $\|\mathbf{x}\|_{\mathbf{A}}^2 = \mathbf{x}^{\mathsf{T}}\mathbf{A}\mathbf{x}$, $\mathbf{L}_+$ is the signless Laplacian matrix of the network, $\mathbf{L}_-$ is the signed Laplacian matrix, $D = \frac{\sigma_n^2(k+1)}{\sigma_n^2(k)}$ is the noise decay factor with $0 < D < 1$, and $\phi_{\max}(\mathbf{A})$ and $\phi_{\min}(\mathbf{A})$ are the non-zero largest and smallest singular values of matrix $\mathbf{A}$, respectively. Hence from (18), the Algorithm 1 converges in mean to the maximum value a^*.

IV. Simulation Results

For numerical results, we consider a network with $L = 5$ agents and edge set $\mathcal{E} = \{(1,2),(1,5),(2,3),(2,4),(3,4),(4,5)\}$. The local state values are chosen as $\mathbf{a} = [a_1, \ldots, a_L]^{\mathsf{T}} = [-2, -1, 0, 1, 2]^{\mathsf{T}}$ and index of node with maximum value $a^* = 2$ is $i = 5$. For the ADMM penalty parameters we chose $\rho_y = \rho_z = 2$. The consensus accuracy is defined as $\epsilon_i(k) = \frac{1}{L}\sum_{i=1}^{L}\frac{|x_i(k)-a^*|^2}{(a^*)^2}$. The perturbation noise $n_i(k)$ is chosen $n_i(k) \sim \mathcal{N}(0, \sigma_n^2(1)/k)$, i.e., $\sigma_n^2(k) = \sigma_n^2(1)/k$. We set initial point $x_i(0) \sim \mathcal{N}(0, \sigma_x^2)$ with variance $\sigma_x^2 = \sigma_n^2(1)$.

Figure 1a illustrates the convergence of the proposed algorithm with $\sigma_x^2 = \sigma_n^2(1) = 10$. We can see that local updates converge to the maximum value at same time guaranteeing privacy. Figure 1b shows the privacy value defined as the estimation error covariance $Q_i(k)$ at the adversary for $i = 1, 2, \ldots, L$. It can be seen that node $i = 5$ has lower privacy compared to other agents since it has a higher probability of the event $\gamma_L(k) = 1$ and thus leading better estimation error at the adversary. Further, the plot shows that privacy depends on the number of neighbors, and agents with same number of neighbors have similar privacy leakage. This follows from (14), where the privacy depends on the measurement noise covariance, which is influenced by the number of neighbours. Here Agents 1 and 3, with $|\mathcal{N}_1| = |\mathcal{N}_3| = 2$, and Agents 2 and 4, with $|\mathcal{N}_2| = |\mathcal{N}_4| = 3$, have same privacy guarantees. The trade-off between privacy and consensus accuracy is shown in Fig. 2. It can be seen that larger privacy leads to lower accuracy. But, the accuracy and privacy can be controlled through appropriate selection of stopping time K and the initial perturbation noise covariance $\sigma_n^2(1)$.

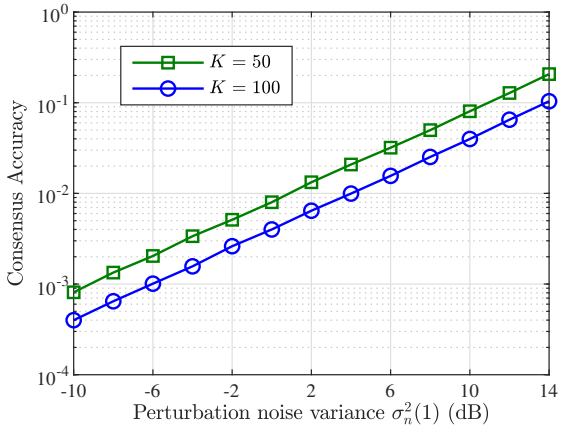

Fig. 2. Accuracy as a function perturbation noise variance $\sigma_n^2(1)$ for $K = 100$ and $K = 50$ ADMM iterations.

V. Conclusion

We have developed a privacy-preserving distributed maximum consensus algorithm where the local state of the agents and identity of the agent with maximum state is kept private from the adversary. We showed that the maximum consensus problem can be recast as a linear program, which is solved in a privacy preserving manner using ADMM and perturbing the primal update step with additive Gaussian noise of decreasing variance. Considering the privacy metric as the estimation error of the local state, we have obtained theoretical bounds on the privacy leakage at the agents.

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] Y. Zhang and S. Li, "Distributed biased min-consensus with applications to shortest path planning," *IEEE Transactions on Automatic Control*, vol. 62, no. 10, pp. 5429–5436, Oct. 2017.

[3] B. Kailkhura, S. Brahma, and P. K. Varshney, "Data falsification attacks on consensus-based detection systems," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 3, no. 1, pp. 145–158, Mar. 2017.

[4] S. Yang, S. Tan, and J. Xu, "Consensus based approach for economic dispatch problem in a smart grid," *IEEE Transactions on Power Systems*, vol. 28, no. 4, pp. 4416–4426, Nov. 2013.

[5] J. Lin, A. S. Morse, and B. D. O. Anderson, "The multi-agent rendezvous problem. part 1: The synchronous case," *SIAM Journal on Control and Optimization*, vol. 46, no. 6, pp. 2096–2119, 2007.

[6] Y. Hong, J. Hu, and L. Gao, "Tracking control for multi-agent consensus with an active leader and variable topology," *Automatica*, vol. 42, no. 7, pp. 1177 – 1182, 2006.

[7] W. B. Heinzelman, A. P. Chandrakasan, and H. Balakrishnan, "An application-specific protocol architecture for wireless microsensor networks," *IEEE Transactions on Wireless Communications*, vol. 1, no. 4, pp. 660–670, Oct. 2002.

[8] E. Nozari, P. Tallapragada, and J. Cortés, "Differentially private average consensus: Obstructions, trade-offs, and optimal algorithm design," *Automatica*, vol. 81, pp. 221–231, Jul. 2017.

[9] J. He, L. Cai, C. Zhao, P. Cheng, and X. Guan, "Privacy-preserving average consensus: Privacy analysis and algorithm design," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 5, no. 1, pp. 127–138, Mar. 2019.

[10] Y. Mo and R. M. Murray, "Privacy preserving average consensus," *IEEE Transactions on Automatic Control*, vol. 62, no. 2, pp. 753–765, Feb. 2017.

[11] Q. Li, I. Cascudo, and M. G. Christensen, "Privacy-preserving distributed average consensus based on additive secret sharing," in *Proc. 27th European Signal Processing Conference*, 2019, pp. 1–5.

[12] R. Lazzeretti, S. Horn, P. Braca, and P. Willett, "Secure multi-party consensus gossip algorithms," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, 2014, pp. 7406–7410.

[13] M. Ambrosin, P. Braca, M. Conti, and R. Lazzeretti, "ODIN: Obfuscation-based privacy-preserving consensus algorithm for decentralized information fusion in smart device networks," *ACM Transactions on Internet Technology*, vol. 18, no. 1, pp. 1–22, Oct. 2017.

[14] M. Ruan, H. Gao, and Y. Wang, "Secure and privacy-preserving consensus," *IEEE Transactions on Automatic Control*, vol. 64, no. 10, pp. 4035–4049, Oct. 2019.

[15] I. D. Ridgley, R. A. Freeman, and K. M. Lynch, "Simple, private, and accurate distributed averaging," in *Proc. 57th Annual Allerton Conference on Communication, Control, and Computing*, 2019, pp. 446–452.

[16] G. Muniraju, C. Tepedelenlioglu, and A. Spanias, "Analysis and design of robust max consensus for wireless sensor networks," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 5, no. 4, pp. 779–791, Dec. 2019.

[17] A. Nowzari and M. G. Rabbat, "Improved bounds for max consensus in wireless networks," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 5, no. 2, pp. 305–319, Jun. 2019.

[18] F. Iutzeler, P. Ciblat, and J. Jakubowicz, "Analysis of max-consensus algorithms in wireless channels," *IEEE Transactions on Signal Processing*, vol. 60, no. 11, pp. 6103–6107, Nov. 2012.

[19] X. Wang, J. He, P. Cheng, and J. Chen, "Differentially private maximum consensus: Design, analysis and impossibility result," *IEEE Transactions on Network Science and Engineering*, vol. 6, no. 4, pp. 928–939, Oct. 2019.

[20] X. Duan, J. He, P. Cheng, Y. Mo, and J. Chen, "Privacy preserving maximum consensus," in *Proc. 54th IEEE Conference on Decision and Control*, 2015, pp. 4517–4522.

[21] E. Nozari, P. Tallapragada, and J. Cortés, "Differentially private distributed convex optimization via functional perturbation," *IEEE Transactions on Control of Network Systems*, vol. 5, no. 1, pp. 395–408, Mar. 2018.

[22] S. Han, U. Topcu, and G. J. Pappas, "Differentially private distributed constrained optimization," *IEEE Transactions on Automatic Control*, vol. 62, no. 1, pp. 50–64, Jan. 2017.

[23] T. Zhang and Q. Zhu, "Dynamic differential privacy for ADMM-based distributed classification learning," *IEEE Transactions on Information Forensics and Security*, vol. 12, no. 1, pp. 172–187, Jan. 2017.

[24] X. Zhang, M. M. Khalili, and M. Liu, "Recycled ADMM: Improving the privacy and accuracy of distributed algorithms," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 1723–1734, 2020.

[25] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "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, Jan. 2010.

[26] G. B. Giannakis, Q. Ling, G. Mateos, and I. D. Schizas, *Splitting Methods in Communication, Imaging, Science, and Engineering*, ser. Scientific Computation, R. Glowinski, S. J. Osher, and W. Yin, Eds. Springer International Publishing, 2016.

[27] P. Braca, R. Lazzeretti, S. Marano, and V. Matta, "Learning with privacy in consensus + obfuscation," *IEEE Signal Processing Letters*, vol. 23, no. 9, pp. 1174–1178, Sep. 2016.

[28] D. Han, K. You, L. Xie, J. Wu, and L. Shi, "Optimal parameter estimation under controlled communication over sensor networks," *IEEE Transactions on Signal Processing*, vol. 63, no. 24, pp. 6473–6485, Dec. 2015.

[29] D. Han, Y. Mo, J. Wu, S. Weerakkody, B. Sinopoli, and L. Shi, "Stochastic event-triggered sensor schedule for remote state estimation," *IEEE Transactions on Automatic Control*, vol. 60, no. 10, pp. 2661–2675, Oct. 2015.

[30] J. Ding, Y. Gong, M. Pan, and Z. Han, "Optimal differentially private ADMM for distributed machine learning," 2019. [Online]. Available: http://arxiv.org/abs/1901.02094