

Preserving Statistical Privacy in Distributed Optimization

Nirupam Gupta[†], Shripad Gade^{*}, Nikhil Chopra[‡] and Nitin H. Vaidya[†]

Abstract—We present a distributed optimization protocol that preserves statistical privacy of agents’ local cost functions against a passive adversary that corrupts some agents in the network. The protocol is a composition of a distributed “zero-sum” obfuscation protocol that obfuscates the agents’ local cost functions, and a standard non-private distributed optimization method. We show that our protocol protects the statistical privacy of the agents’ local cost functions against a passive adversary that corrupts up to t arbitrary agents as long as the communication network has $(t + 1)$ -vertex connectivity. The “zero-sum” obfuscation protocol preserves the sum of the agents’ local cost functions and therefore ensures accuracy of the computed solution.

Index Terms—Statistical privacy, Distributed optimization, Large-scale systems, Sensor networks.

I. INTRODUCTION

Distributed optimization in multi-agent peer-to-peer networks has gained significant attention in recent years [1]. In this problem, each agent has a local cost function and the goal for the agents is to collectively minimize sum of their local cost functions. Specifically, we consider n agents, where each agent i has a convex cost $h_i : \mathbb{R}^m \rightarrow \mathbb{R}$ and a convex, compact set \mathcal{X} . A *distributed optimization algorithm* enables the agents to collectively compute a *global* minimum,

$$x^* \in \arg \min_{x \in \mathcal{X}} \sum_{i=1}^n h_i(x). \quad (1)$$

We consider a scenario when a *passive* adversary can corrupt some of the agents in the network. The corrupted agents follow the prescribed protocol correctly, but may try to learn about the cost functions of other non-corrupted agents in the network. In literature, a passive adversary is also commonly referred as *honest-but-curious*. Prior work has shown that for certain distributed optimization algorithms, such as the Distributed Gradient Descent (DGD) method, a passive adversary may learn about all the agents’ cost functions by corrupting only a subset of agents in the network [2]. This is clearly undesirable in general, and especially in cases where the cost functions may contain sensitive information [3].

In this paper, we consider the Function Sharing (FS) protocol [4], wherein the agents obfuscate their local cost functions with *correlated random functions* before executing a (non-private) distributed optimization algorithm such as the DGD

method. The obfuscation strategy is aggregate invariant by construction and therefore, the agents compute a minimizer (1) accurately using solely their obfuscated local cost functions [2, Theorem 1]. The FS protocol was first proposed by Gade et al. [2]. However, as of yet, the FS protocol lacks a formal privacy analysis. In this paper, we utilize the statistical privacy definition developed by Gupta et al. [5], [6] to present a privacy guarantee of the FS protocol.

In the past, distributed optimization protocols have been proposed for preserving differential privacy of the agents’ local cost functions. However, these differentially private protocols suffer inevitably from privacy-accuracy trade-offs [7], [8]. That is, the agents can only compute an approximation of a global minimum x^* , defined by (1). The FS protocol allows the agents to compute a global minimum (1) accurately, and therefore, it obtains a weaker statistical privacy guarantee compared to the differentially private protocols.

Homomorphic encryption-based privacy protocols implicitly rely on two pragmatic assumptions, (1) computational intractability of *hard* mathematical problems, and (2) limited computational power of a passive adversary [3], [9]–[11]. We show that the FS protocol provides *statistical* (or information-theoretic [12]) privacy, which is valid regardless of the above assumptions.

However, both the differential privacy based protocols and the homomorphic encryption based protocols can provide privacy against eavesdroppers [3], [7]–[11]. The FS protocol, on the other hand, can only provide privacy against honest-but-curious agents in the network.

Summary of Our Contributions: We show that in the FS protocol the passive adversary obtains limited information, in a statistical sense, about the local cost functions of the non-corrupted (or *honest*) agents, as long as the agents corrupted by the passive adversary do not form a *vertex cut* in the underlying communication network topology. Thus, the FS protocol protects the statistical privacy of the honest agents’ local cost functions against any passive adversary that corrupts up to t arbitrary agents in the system as long as the communication network topology has $(t + 1)$ -vertex connectivity.

It is of independent interest to note that a variant of the FS protocol is known to preserve the *perfect* statistical privacy in distributed average consensus problem [13]–[15].

II. PROBLEM SETUP

We consider a passive adversary, denoted by \mathcal{A} , that corrupts some agents in the network. The goal is to design distributed optimization protocols that protect the privacy of the non-corrupted (or *honest*) agents’ local cost functions against

Partially supported by NSF through grant ECCS1711554, award 1842198, and ARL under Cooperative Agreement W911NF-17-2-0196.

[†] Department of Computer Science, Georgetown University, Washington, DC 20057, USA. ^{*} Electrical and Computer Engineering, University of Illinois Urbana-Champaign, Urbana, IL, USA. [‡] Mechanical Engineering, University of Maryland, College Park, MD, USA. Email: nirupam115@gmail.com, gade3@illinois.edu, nchopra@umd.edu, and nitin.vaidya@georgetown.edu.

the passive adversary, while allowing the agents to compute solution (1) accurately. The adversary is assumed passive and the corrupted agents execute the prescribed protocol correctly. For a distributed optimization protocol Π , we define *view* of \mathcal{A} for an execution of Π as follows.

Definition 1. For a protocol Π , the *view* of \mathcal{A} constitutes the information stored, transmitted and received by the agents corrupted by \mathcal{A} during the execution of Π .

Privacy requires that the entire *view* of \mathcal{A} does not leak significant (or any) information about the local costs of the honest agents. Note that, by definition, \mathcal{A} inevitably learns a point $x^* \in \arg \min_{x \in \mathcal{X}} \sum_{i=1}^n h_i(x)$, assuming it corrupts at least one agent. A *perfectly* private protocol would not reveal any information about the honest agents' cost functions to \mathcal{A} besides x^* . However, such a perfect privacy is quite difficult to guarantee. For now, we relax the privacy requirement, and only consider privacy for the *affine* terms of the agents' cost functions. However, as elaborated in Section III-B, the FS protocol can be extended easily for privacy of *higher-order polynomial* terms. That is, we implicitly assume that the non-affine terms of the agents' cost functions are known a priori to the passive adversary.

For each agent i , the cost function $h_i(x)$ can be decomposed into two parts; the *affine* term denoted by $h_i^{(1)}(x)$, and the *non-affine* term denoted by $h_i^\dagger(x)$. Specifically,

$$h_i(x) = h_i^{(1)}(x) + h_i^\dagger(x), \quad \forall x \in \mathbb{R}^m, i \in \{1, \dots, n\}. \quad (2)$$

As the name suggests, the affine terms are affine in x . That is, for each i there exists $\alpha_i \in \mathbb{R}^m$ and $\gamma_i \in \mathbb{R}$ such that, $h_i^{(1)}(x) = \alpha_i^T x + \gamma_i$, $\forall x \in \mathbb{R}^m$, where $(\cdot)^T$ denotes the transpose. As the constants γ_i 's do not affect the solution of the optimization problem (1), the agents need not share these constants with each other. Hence, the privacy of honest agents' γ_i 's can be trivially preserved. For a meaningful discussion of privacy we will ignore these constants. Let,

$$A = [\alpha_1, \dots, \alpha_n] \quad (3)$$

be the $m \times n$ -dimensional matrix obtained by column-wise stacking of the individual agents' affine coefficients.

Let \mathcal{C} denote the set of agents corrupted by the adversary \mathcal{A} , and let \mathcal{H} denote the remaining non-corrupted (or honest) agents. For privacy preservation, the protocol Π may introduce some randomness in the system, in which case the *view* of \mathcal{A} is a random variable. Let,

- $\text{View}_{\mathcal{C}}(A)$ denote the probability distribution of the *view* of \mathcal{A} for an execution of Π when the agents' private cost functions have affine coefficients A .

Our definition of privacy below is built on *relative entropy*, which is also known as the Kullback-Leibler (KL) divergence. For a continuous random variable \mathcal{R} , let $f_{\mathcal{R}}(r)$ denote its probability distribution or probability density function (p.d.f.) at $r \in \mathcal{R}$. The KL-divergence, denoted by D_{KL} , quantifies the difference between a certain probability distribution $f'_{\mathcal{R}}$ and the reference probability distribution $f_{\mathcal{R}}$ [16]. Specifically, the KL-divergence of $f'_{\mathcal{R}}$ from $f_{\mathcal{R}}$ is defined as

$$D_{KL}(f_{\mathcal{R}}, f'_{\mathcal{R}}) = \int_{\mathcal{R}} f_{\mathcal{R}}(s) \log \left(\frac{f_{\mathcal{R}}(s)}{f'_{\mathcal{R}}(s)} \right) ds.$$

Let $\|\cdot\|$ denote the Euclidean norm for vectors and the Frobenius norm for matrices.

Definition 2. For $\epsilon > 0$, a distributed optimization protocol Π is said to be " (\mathcal{C}, ϵ) -affine private" if for every pair of agents' affine coefficients $A = [\alpha_1, \dots, \alpha_n]$ and $B = [\beta_1, \dots, \beta_n]$ subject to the constraints:

$$\alpha_i = \beta_i, \quad \forall i \in \mathcal{C}, \quad \text{and} \quad \sum_{i \in \mathcal{H}} \alpha_i = \sum_{i \in \mathcal{H}} \beta_i, \quad (4)$$

the supports of $\text{View}_{\mathcal{C}}(B)$ and $\text{View}_{\mathcal{C}}(A)$ are identical, and

$$D_{KL}(\text{View}_{\mathcal{C}}(A), \text{View}_{\mathcal{C}}(B)) \leq \epsilon \|A - B\|^2. \quad (5)$$

In other words, Definition 2 implies that if Π is (\mathcal{C}, ϵ) -affine private then an adversary \mathcal{A} cannot unambiguously distinguish between two sets of agents' affine coefficients, A and B , that are identical for the corrupted agents and have identical sum over all honest agents (i.e., satisfy (4)). The value of ϵ signifies the *strength* of the privacy obtained. Smaller is the value of ϵ , the more difficult it is for \mathcal{A} to distinguish between two sets of agents' affine coefficients satisfying (4), and hence stronger is the privacy.

III. PROPOSED PROTOCOL AND PRIVACY GUARANTEE

In this section, we present the Function Sharing (FS) protocol and the formal privacy guarantee.

The notation used is as follows. The underlying communication network is modeled by an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where the set of nodes $\mathcal{V} = \{1, \dots, n\}$ denotes the agents (indexed arbitrarily), and the set of edges \mathcal{E} denotes the communication links between the agents. Being undirected, each edge $e \in \mathcal{E}$ is represented by an unordered pair of agents. For each i , the set $\mathcal{N}_i = \{j \in \mathcal{V} \mid \{i, j\} \in \mathcal{E}\}$ denotes the neighbors of agent i .

The FS protocol constitutes two phases as elaborated in Algorithm 1. In phase I, each agent i uses a "zero-sum" obfuscation protocol to compute an "effective cost function" $\tilde{h}_i(x)$ based on its private local cost function $h_i(x)$. In phase II, the agents use the DGD algorithm on their effective local cost functions to solve for the effective optimization problem,

$$\underset{x \in \mathcal{X}}{\text{minimize}} \sum_{i=1}^n \tilde{h}_i(x). \quad (6)$$

We now show that upon completion of phase II the agents indeed obtain a common minimum of the original optimization problem (1). As \mathcal{G} is an undirected graph,

$$\sum_{i=1}^n u_i = \sum_{i=1}^n \sum_{j \in \mathcal{N}_i} (r_{ji} - r_{ij}) = 0.$$

This implies that, for all $x \in \mathbb{R}^m$,

$$\sum_{i=1}^n \tilde{h}_i(x) = \sum_{i=1}^n h_i(x) + \sum_{i=1}^n u_i^T x = \sum_{i=1}^n h_i(x). \quad (7)$$

Equivalently, the masking in phase I preserves the sum of the agents' local cost functions. Therefore, a solution for problem (6), obtained using the DGD algorithm in [17], is a solution for the original optimization problem (1).

Algorithm 1 Function Sharing (FS) Protocol

Input: Each agent i has cost function $h_i(x)$ and $\sigma \in \mathbb{R}$.

Output: Minimizer, $x^* \in \arg \min_{x \in \mathcal{X}} \sum_{i=1}^n h_i(x)$

◇ **Phase 1: Masking of Cost Functions**

Each agent $i \in \mathcal{V}$ executes:

- 1: Draws vectors $r_{ij} \sim N(0_m, \sigma^2 I_m)$ independently for $j \in \mathcal{N}_i$ and sends r_{ij} to each agent $j \in \mathcal{N}_i$.
- 2: Compute the *mask* u_i

$$u_i = \sum_{j \in \mathcal{N}_i} (r_{ij} - r_{ji}) \quad (8)$$

- 3: Compute the *effective cost function* $\tilde{h}_i(x)$,

$$\tilde{h}_i(x) = h_i(x) + u_i^T x, \quad \forall x \in \mathbb{R}^m. \quad (9)$$

◇ **Phase 2: Distributed Optimization**

- 4: Agents execute the DGD algorithm [18] on the local effective costs $\{\tilde{h}_i(x)\}_{i \in \mathcal{V}}$.
-

A. Privacy Guarantee

The privacy guarantee for the above FS protocol is presented by Theorem 1 below. Recall that \mathcal{C} denotes the set of agents corrupted by the passive adversary. Let $\mathcal{H} = \mathcal{V} \setminus \mathcal{C}$ denote the set of honest agents, and let $\mathcal{G}_{\mathcal{H}}$ denote the residual graph obtained by removing the agents in \mathcal{C} , and the edges incident to them, from \mathcal{G} . Let $\mathcal{L}_{\mathcal{H}}$ denote the graph-Laplacian of $\mathcal{G}_{\mathcal{H}}$ and $\underline{\mu}(\mathcal{L}_{\mathcal{H}})$ denote the second smallest eigenvalue of $\mathcal{L}_{\mathcal{H}}$. The eigenvalue $\underline{\mu}(\mathcal{L}_{\mathcal{H}})$ is also commonly known as the *algebraic connectivity* of the graph [19].

Theorem 1. *If \mathcal{C} is not a vertex cut of \mathcal{G} , and the affine coefficients of the agents' private cost functions are independent of each other, then the FS protocol is (\mathcal{C}, ϵ) -affine private, with $\epsilon = 1/(4\sigma^2 \underline{\mu}(\mathcal{L}_{\mathcal{H}}))$.*

Theorem 1 implies that \mathcal{C} not being a vertex cut¹ of \mathcal{G} is sufficient for (\mathcal{C}, ϵ) -affine privacy. Note that, smaller the value of ϵ , stronger is the privacy. According to Theorem 1, ϵ is inversely proportional both to the variance σ^2 of the elements of random vectors r_{ij} 's used for masking of agents' local costs, and the *algebraic connectivity* of the residual network topology $\mathcal{G}_{\mathcal{H}}$. Therefore, the agents can achieve stronger privacy by using random vectors with larger variances (i.e., larger σ^2) in phase I of the FS protocol. Additionally, FS protocol guarantees stronger privacy if the residual honest graph $\mathcal{G}_{\mathcal{H}}$ is densely connected.

We further note that the FS protocol can guarantee privacy against any passive adversary that corrupts at most t agents in the network if the network has $(t+1)$ -vertex connectivity. Specifically, we have the following corollary of Theorem 1.

Corollary 1. *If \mathcal{G} has $(t+1)$ -vertex connectivity and the affine coefficients of the agents' private cost functions are independent of each other, then for an arbitrary set $\mathcal{C} \subseteq \mathcal{V}$*

¹A *vertex cut* is a set of vertices of a graph which, if removed – together with any incident edges – disconnects the graph [19].

with $|\mathcal{C}| \leq t$ the FS protocol is (\mathcal{C}, ϵ) -affine private with

$$\epsilon = \max \left\{ \frac{1}{4\sigma^2 \underline{\mu}(\mathcal{L}_{\mathcal{H}})} \mid \mathcal{H} = \mathcal{V} \setminus \mathcal{C}, |\mathcal{C}| \leq t \right\}.$$

The above connectivity condition for privacy is indeed tight. Specifically, the $(t+1)$ -vertex connectivity is *necessary* for privacy against at most t colluding honest-but-curious agents in the consensus-based distributed gradient and subgradient optimization algorithms [2], [20], [21].

B. Privacy of Higher-Degree Polynomial Terms

The FS protocol presented in Algorithm 1 only protects the privacy of affine coefficients of local cost functions, as formally stated in Theorem 1. In what follows, we show an easy extension to protect privacy of higher degree polynomial terms of agents' private cost functions. Here, we assume the agents' cost functions to be univariate, i.e., $x \in \mathbb{R}$.

For each agent i , let $\alpha_i^{(\ell)}$ denote the ℓ -th degree coefficient of its cost function $h_i(x)$. Similar to the definition of (\mathcal{C}, ϵ) -affine privacy, we now define the privacy of the ℓ -th degree coefficients $A^{(\ell)} = [\alpha_1^{(\ell)}, \dots, \alpha_n^{(\ell)}]$ against a passive adversary that corrupts a set of agents \mathcal{C} . Let $\text{View}_{\mathcal{C}}(A^{(\ell)})$ denote the probability distribution of the *view* of adversary \mathcal{A} when ℓ -th degree coefficients of agents' private cost functions are given by $A^{(\ell)}$.

Privacy Definition: For $\epsilon > 0$, protocol Π is said to preserve the (\mathcal{C}, ϵ) -privacy of ℓ -th degree coefficients $A^{(\ell)}$ if for every other set of ℓ -th degree coefficients $B = [\beta_1^{(\ell)}, \dots, \beta_n^{(\ell)}]$ subject to the constraints:

$$\beta_i^{(\ell)} = \alpha_i^{(\ell)}, \quad \forall i \in \mathcal{C}, \quad \text{and} \quad \sum_{i \in \mathcal{H}} \beta_i^{(\ell)} = \sum_{i \in \mathcal{H}} \alpha_i^{(\ell)},$$

the support of $\text{View}_{\mathcal{C}}(A^{(\ell)})$ & $\text{View}_{\mathcal{C}}(B^{(\ell)})$ are identical, and

$$D_{KL} \left(\text{View}_{\mathcal{C}}(A^{(\ell)}), \text{View}_{\mathcal{C}}(B^{(\ell)}) \right) \leq \epsilon \|A^{(\ell)} - B^{(\ell)}\|^2.$$

When defining the distribution $\text{View}_{\mathcal{C}}(A^{(\ell)})$, we implicitly assume that the passive adversary \mathcal{A} knows all the coefficients of the honest agents' costs, except the ℓ -th coefficients $\{\alpha_i^{(\ell)}, i \in \mathcal{H}\}$. Thus, the privacy analysis here is conservative.

Modified FS Protocol and Privacy Guarantee: In the first phase, the agents mask the coefficients $A^{(\ell)}$ in a similar manner as the masking of the affine coefficients delineated in Algorithm 1 to compute the effective cost functions.

Note that in this case, due to the non-affine masking, the effective cost functions $\tilde{h}_i(x)$'s may become non-convex. The sum of the effective cost functions, however, is still a convex function (see (7)). As discussed in [17], the DGD algorithm allows agents to minimize convex sum of their local non-convex cost functions, provided that the local cost functions' gradients are Lipschitz continuous [2, Theorem 1]. The DGD can be substituted with other distributed optimization algorithms, provided those algorithms also minimize convex sum of non-convex functions (see [17] for details).

Now, Theorem 1 implies that if \mathcal{C} does not form a vertex cut of the network topology \mathcal{G} then the FS protocol, modified as

above, preserves the (\mathcal{C}, ϵ) -privacy of ℓ -th degree coefficients $A^{(\ell)}$ for each $\ell = \{1, \dots, d\}$, where, privacy parameter $\epsilon = 1/(4\sigma^2\mu(\mathcal{L}))$.

IV. PROOF OF THEOREM 1

In this section, we present the formal proof for Theorem 1. In principle, the proof is a generalization of the privacy analysis presented in [5]. First, we state a few critical observations in Lemmas 1 and 2 below.

Let \mathcal{L} denote the graph-Laplacian of the network topology \mathcal{G} . As \mathcal{G} is undirected, \mathcal{L} is a diagonalizable matrix [19]. Specifically, there exists a unitary matrix M constituting the orthogonal eigenvectors of \mathcal{L} such that², $\mathcal{L} = M \text{Diag}(\mu_1, \dots, \mu_n) M^T$ where $\mu_1 \leq \mu_2 \leq \dots \leq \mu_n$ are the eigenvalues of \mathcal{L} . When \mathcal{G} is connected, $\mu_1 = 0$ and $\mu_2 > \mu_1$ [19]. We denote the *generalized inverse* of \mathcal{L} by \mathcal{L}^\dagger . Note that [22],

$$\mathcal{L}^\dagger = M \text{Diag}(0, 1/\mu_2, \dots, 1/\mu_n) M^T \quad (10)$$

For future usage, we denote the second smallest eigenvalue of \mathcal{L} , i.e., μ_2 , by $\underline{\mu}(\mathcal{L})$. Let 0_n and 1_n denote the zero and the one vectors, respectively, of dimension n . For a positive real value c , $N^\dagger(0_n, c\mathcal{L})$ denotes the *degenerate* Gaussian distribution [23]. Specifically, if $\mathcal{R} \sim N^\dagger(0_n, c\mathcal{L})$ and \mathcal{G} is a connected graph then,

$$f_{\mathcal{R}}(r) = \begin{cases} \frac{1}{\sqrt{\det^*(2\pi c\mathcal{L})}} \exp\left(-\frac{r^T \mathcal{L}^\dagger r}{2c}\right) & , r^T 1_n = 0 \\ 0 & , \text{otherwise} \end{cases} \quad (11)$$

where $\det^*(2\pi c\mathcal{L}) = (2\pi c)^{n-1} \prod_{i=2}^n \mu_i$. Henceforth, for a vector v , v^k denotes the k -th element of v unless otherwise noted. For $i \in \mathcal{V}$, recall that u_i is the *mask* (see (8)). Let,

$$U^k = [u_1^k, \dots, u_n^k]^T, \quad k = 1, \dots, m. \quad (12)$$

be a n -dimensional vector comprising the k -th elements of the *masks* computed by the agents in phase I of the FS protocol. For a random vector \mathcal{R} , we denote its mean by $\mathbb{E}(\mathcal{R})$ and its covariance matrix by $\text{Cov}(\mathcal{R})$. Note that $\text{Cov}(\mathcal{R}) = \mathbb{E}(\mathcal{R} - \mathbb{E}(\mathcal{R}))(\mathcal{R} - \mathbb{E}(\mathcal{R}))^T$.

Lemma 1. *If \mathcal{G} is a connected graph then for each $k \in \{1, \dots, m\}$, $U^k \sim N^\dagger(0_n, 2\sigma^2\mathcal{L})$.*

Proof: Assign an arbitrary order to the set of edges, i.e., let $\mathcal{E} = (e_1, \dots, e_{|\mathcal{E}|})$. For each edge e_l where $l \in \{1, \dots, |\mathcal{E}|\}$, we define a vector θ_{e_l} of size n whose i -th element denoted by $\theta_{e_l}^i$ is given as follows:

$$\theta_{e_l}^i = \begin{cases} 1 & \text{if } e_l = \{i, j\} \text{ and } i < j \\ -1 & \text{if } e_l = \{i, j\} \text{ and } i > j \\ 0 & \text{otherwise.} \end{cases}$$

Let $\Theta = [\theta_{e_1}, \dots, \theta_{e_{|\mathcal{E}|}}]$ be an *oriented incidence matrix* of graph \mathcal{G} [19]. For each edge $e = \{i, j\}$ with $i < j$,

$$c_e \triangleq r_{ji} - r_{ij}. \quad (13)$$

Since the each random vector in $\{r_{ij}, i, j \in \mathcal{V}\}$ is identically and independently distributed (i.i.d.) by a normal distribution

² $\text{Diag}(y_1, \dots, y_n)$ is a diagonal matrix with diagonal entries y_1, \dots, y_n .

$N(0_m, \sigma^2 I_m)$, (13) implies that for each edge e_l the random vector c_{e_l} is i.i.d. as $N(0, 2\sigma^2 I_m)$. Therefore, for each k , the random variable $c_{e_l}^k$ has normal distribution of $N(0, 2\sigma^2)$. Let, $C^k = [c_{e_1}^k, c_{e_2}^k, \dots, c_{e_{|\mathcal{E}|}}^k]^T$. For two distinct edges e and e' , the random vectors c_e and $c_{e'}$ are independent. Therefore,

$$\mathbb{E}(C^k)(C^k)^T = 2\sigma^2 I_{|\mathcal{E}|}, \quad (14)$$

where $I_{|\mathcal{E}|}$ is $|\mathcal{E}| \times |\mathcal{E}|$ identity matrix. Moreover, from (8), $U^k = \Theta C^k$, $\forall k \in \{1, \dots, m\}$. As \mathcal{G} is assumed connected, the support of U^k is the entire space orthogonal to 1_n . Also, $\mathbb{E}(U^k) = \Theta \mathbb{E}(C^k) = 0_n$. As $\mathcal{L} = \Theta \Theta^T$ [19], $\text{Cov}(U^k) = \Theta (\mathbb{E}(C^k)(C^k)^T) \Theta^T = 2\sigma^2 \Theta \Theta^T = 2\sigma^2 \mathcal{L}$. Thus, U^k has the generalized Gaussian distribution $N^\dagger(0_n, 2\sigma^2 \mathcal{L})$. ■

Using the above lemma, we show that the knowledge of the *effective cost functions* does not provide significant information about the *affine coefficients* of the agents' private cost functions.

Consider two possible executions E_A and E_B of the FS protocol such that the affine coefficients of the agents' effective cost functions in both executions are given by $\tilde{A} = [\tilde{\alpha}_1, \dots, \tilde{\alpha}_n]$. In execution E_A , the agents have local cost functions with affine coefficients $A = [\alpha_1, \dots, \alpha_n]$, and in execution E_B , the agents have local cost functions with affine coefficients $B = [\beta_1, \dots, \beta_n]$. Let $f_{\tilde{A}|A}$ and $f_{\tilde{A}|B}$ denote the conditional p.d.f.s of \tilde{A} given that the affine coefficients of the agents' private cost functions are A and B , respectively. Recall that $\underline{\mu}(\mathcal{L})$ denotes the second smallest eigenvalue of the graph-Laplacian matrix \mathcal{L} , i.e., μ_2 .

Lemma 2. *If \mathcal{G} is connected, and $\sum_{i=1}^n \alpha_i = \sum_{i=1}^n \beta_i$, then supports of $f_{\tilde{A}|A}$ and $f_{\tilde{A}|B}$ are identical, and*

$$D_{KL}(f_{\tilde{A}|A}, f_{\tilde{A}|B}) \leq \frac{1}{4\sigma^2 \underline{\mu}(\mathcal{L})} \|A - B\|^2. \quad (15)$$

Proof: Let, \tilde{A}^k and A^k denote the column vectors representing the k -th rows of the effective affine coefficients \tilde{A} and the actual affine coefficients A , respectively. That is, $\tilde{A}^k = [\tilde{\alpha}_1^k, \dots, \tilde{\alpha}_n^k]^T$ and $A^k = [\alpha_1^k, \dots, \alpha_n^k]^T$. The proof comprises three parts.

Part I: Recall from (8), $\tilde{\alpha}_i^k = \alpha_i^k + u_i^k$ for all i and k . Therefore (see (12) for the notation U^k), $\tilde{A}^k = A^k + U^k$. As U^k is independent of A^k for every k , we get,

$$f_{\tilde{A}^k|A^k}(\tilde{\alpha}_1^k, \dots, \tilde{\alpha}_n^k) = f_{U^k}(\tilde{A}^k - A^k). \quad (16)$$

Therefore, from Lemma 1, if $\sum_{i=1}^n \tilde{\alpha}_i^k = \sum_{i=1}^n \alpha_i^k$ then,

$$f_{\tilde{A}^k|A^k}(\tilde{\alpha}_1^k, \dots, \tilde{\alpha}_n^k) = \frac{1}{\sqrt{\det^*(4\pi\sigma^2\mathcal{L})}} \exp\left(-\frac{(\tilde{A}^k - A^k)^T \mathcal{L}^\dagger (\tilde{A}^k - A^k)}{4\sigma^2}\right) \quad (17)$$

Else if $\sum_{i=1}^n \tilde{\alpha}_i^k \neq \sum_{i=1}^n \alpha_i^k$ then

$$f_{\tilde{A}^k|A^k}(\tilde{\alpha}_1^k, \dots, \tilde{\alpha}_n^k) = 0, \quad (18)$$

From (17) and (18), it is easy to see that the supports of the conditional p.d.f.s $f_{\tilde{A}^k|A}$ and $f_{\tilde{A}^k|B}$ are identical.

Part II: From (17),

$$\log \frac{f_{\tilde{A}^k|A^k}(\tilde{\alpha}_1^k, \dots, \tilde{\alpha}_n^k)}{f_{\tilde{A}^k|B^k}(\tilde{\alpha}_1^k, \dots, \tilde{\alpha}_n^k)} = \frac{(A^k - B^k)^T \mathcal{L}^\dagger (2\tilde{A}^k - A^k - B^k)}{4\sigma^2}$$

Let $s = \tilde{A}^k - A^k$, then we get, $D_{KL}(f_{\tilde{A}^k|A^k}, f_{\tilde{A}^k|B^k}) =$

$$\begin{aligned} & \frac{1}{4\sigma^2} \int_{s \in \mathbb{R}^n} (A^k - B^k)^T \mathcal{L}^\dagger (2s + A^k - B^k) f_{U^k}(s) ds \\ & = \frac{1}{2\sigma^2} (A^k - B^k)^T \mathcal{L}^\dagger \mathbb{E}(U^k) + \frac{1}{4\sigma^2} (A^k - B^k)^T \mathcal{L}^\dagger (A^k - B^k). \end{aligned}$$

From Lemma 1, $\mathbb{E}(U^k) = 0_n$. Therefore,

$$D_{KL}(f_{\tilde{A}^k|A^k}, f_{\tilde{A}^k|B^k}) = \frac{1}{4\sigma^2} (A^k - B^k)^T \mathcal{L}^\dagger (A^k - B^k). \quad (19)$$

As \mathcal{G} is assumed connected, $\text{rank}(\mathcal{L}) = n - 1$ and $\mathcal{L}1_n = 0_n$. Recall that $1_n^T (A^k - B^k) = 0_n$. Thus, the vector $A^k - B^k$ belongs to the space orthogonal to the nullspace of \mathcal{L} . Now, substituting \mathcal{L}^\dagger from (10) in (19) we obtain that

$$D_{KL}(f_{\tilde{A}^k|A^k}, f_{\tilde{A}^k|B^k}) \leq \frac{\|A^k - B^k\|^2}{4\sigma^2 \underline{\mu}(\mathcal{L})}. \quad (20)$$

Part III: For $k \neq l$, U^k, U^l are independent of each other. From (16), $f_{\tilde{A}^k|A^k} = \prod_{k=1}^m f_{\tilde{A}^k|A^k}$, and similarly, $f_{\tilde{A}^k|B^k} = \prod_{k=1}^m f_{\tilde{A}^k|B^k}$. This, due to the KL-divergence property, implies that

$$D_{KL}(f_{\tilde{A}|A}, f_{\tilde{A}|B}) = \sum_{k=1}^m D_{KL}(f_{\tilde{A}^k|A^k}, f_{\tilde{A}^k|B^k}).$$

Substituting from (20) above concludes the proof. \blacksquare

Theorem 1 can be now proved easily using Lemma 2.

Proof of Theorem 1. Recall that \mathcal{C} denotes the set of corrupted agents and $\mathcal{H} = \mathcal{V} \setminus \mathcal{C}$ denotes the set of honest agents. Let $\mathcal{E}_{\mathcal{C}}$ denote set of edges incident to \mathcal{C} and $\mathcal{E}_{\mathcal{H}} = \mathcal{E} \setminus \mathcal{E}_{\mathcal{C}}$ be the set of edges incident only to honest agents.

Let the agents' true affine coefficients be given by an $m \times n$ -dimensional matrix $A = [\alpha_1, \dots, \alpha_n]$, as defined in (3). Recall the definition of $\text{View}_{\mathcal{C}}(A)$ from Section II. In this part, we derive the p.d.f. of $\text{View}_{\mathcal{C}}(A)$ for the FS protocol, assuming the worst-case scenario where the *effective cost functions* of all the agents are revealed to the corrupted agents in the second phase. From Definition 1, note that the *view* of the adversary \mathcal{A} for the FS protocol comprises the following information:

- 1) The corrupted agents' private and *effective* cost functions, i.e., $\{h_i(x), \tilde{h}_i(x), i \in \mathcal{C}\}$.
- 2) The set of random vectors $R_{\mathcal{C}} = \{r_{ij}, \{i, j\} \in \mathcal{E}_{\mathcal{C}}\}$.
- 3) The *effective cost functions* of the honest agents, i.e., $\{\tilde{h}_i(x), i \in \mathcal{H}\}$.

For each agent $i \in \mathcal{V}$, let $\tilde{\alpha}_i$ denote the affine coefficient of $\tilde{h}_i(x)$. Let $\tilde{A}_{\mathcal{C}} = [\tilde{\alpha}_i, i \in \mathcal{C}]$ and $\tilde{A}_{\mathcal{H}} = [\tilde{\alpha}_i, i \in \mathcal{H}]$ be the collection of the effective affine coefficients of the corrupted and the honest agents, respectively. Let $f_{(\tilde{A}_{\mathcal{H}}, \tilde{A}_{\mathcal{C}}, R_{\mathcal{C}})|A}$ denote the conditional joint p.d.f. of $\tilde{A}_{\mathcal{H}}, \tilde{A}_{\mathcal{C}}$ and $R_{\mathcal{C}}$ given the agents' true affine coefficients A . From above we obtain that

$$\text{View}_{\mathcal{C}}(A) = f_{(\tilde{A}_{\mathcal{H}}, \tilde{A}_{\mathcal{C}}, R_{\mathcal{C}})|A}. \quad (21)$$

For each agent $i \in \mathcal{V}$, let $\mathcal{C}_i = \mathcal{N}_i \cap \mathcal{C}$. Note that, see (9),

$$\tilde{\alpha}_i = \alpha_i + \sum_{j \in \mathcal{N}_i \setminus \mathcal{C}_i} (r_{ij} - r_{ji}) + \sum_{j \in \mathcal{C}_i} (r_{ij} - r_{ji}), \quad \forall i. \quad (22)$$

For each honest agent $i \in \mathcal{H}$, let

$$\bar{\alpha}_i = \alpha_i + \sum_{j \in \mathcal{N}_i \setminus \mathcal{C}_i} (r_{ij} - r_{ji}). \quad (23)$$

Let $\bar{A}_{\mathcal{H}} = [\bar{\alpha}_i, i \in \mathcal{H}]$ be the collection of honest agents' $\bar{\alpha}_i$'s. Recall that, for two agents i and j , the vectors $r_{ij}, r_{ji} \in R_{\mathcal{C}}$ if and only if $i \in \mathcal{C}$ or $j \in \mathcal{C}$. Therefore, for each honest agent $i \in \mathcal{H}$, the value of $\sum_{j \in \mathcal{C}_i} (r_{ij} - r_{ji})$ is deterministic given $R_{\mathcal{C}}$. Thus,

$$f_{(\tilde{A}_{\mathcal{H}}, \tilde{A}_{\mathcal{C}}, R_{\mathcal{C}})|A} = f_{(\bar{A}_{\mathcal{H}}, \tilde{A}_{\mathcal{C}}, R_{\mathcal{C}})|A}. \quad (24)$$

As the agents' affine coefficients are assumed independent of each other, we have from (23), $\bar{A}_{\mathcal{H}}$ is independent of $\tilde{A}_{\mathcal{C}}$. Moreover, (23) also implies that $\bar{A}_{\mathcal{H}}$ is independent of $R_{\mathcal{C}}$. Therefore, $f_{(\bar{A}_{\mathcal{H}}, \tilde{A}_{\mathcal{C}}, R_{\mathcal{C}})|A} = f_{\bar{A}_{\mathcal{H}}|A} f_{(\tilde{A}_{\mathcal{C}}, R_{\mathcal{C}})|A}$. Note that (i) $\tilde{A}_{\mathcal{C}}$ and $R_{\mathcal{C}}$ are independent of the honest agents' affine coefficients $A_{\mathcal{H}} = [\alpha_i, i \in \mathcal{H}]$, and (ii) $\bar{A}_{\mathcal{H}}$ is also independent of the corrupted agents' affine coefficients $A_{\mathcal{C}} = [\alpha_i, i \in \mathcal{C}]$. Thus, $f_{(\bar{A}_{\mathcal{H}}, \tilde{A}_{\mathcal{C}}, R_{\mathcal{C}})|A} = f_{\bar{A}_{\mathcal{H}}|A_{\mathcal{H}}} f_{(\tilde{A}_{\mathcal{C}}, R_{\mathcal{C}})|A_{\mathcal{C}}}$. Upon substituting this in (24), and using (21), we obtain that

$$\text{View}_{\mathcal{C}}(A) = f_{\bar{A}_{\mathcal{H}}|A_{\mathcal{H}}} f_{(\tilde{A}_{\mathcal{C}}, R_{\mathcal{C}})|A_{\mathcal{C}}}. \quad (25)$$

Now, consider an alternate scenario where the agents' collective affine coefficients are $B = [\beta_1, \dots, \beta_n]$, such that $\beta_i = \alpha_i, \forall i \in \mathcal{C}$, and $\sum_{i \in \mathcal{V}} \beta_i = \sum_{i \in \mathcal{V}} \alpha_i$. Using similar arguments as above, we will obtain that

$$\text{View}_{\mathcal{C}}(B) = f_{\bar{B}_{\mathcal{H}}|B_{\mathcal{H}}} f_{(\tilde{B}_{\mathcal{C}}, R_{\mathcal{C}})|B_{\mathcal{C}}} \quad (26)$$

where $\bar{B}_{\mathcal{H}}, B_{\mathcal{H}}, \tilde{B}_{\mathcal{C}}$ and $B_{\mathcal{C}}$ are the counterparts of $\bar{A}_{\mathcal{H}}, A_{\mathcal{H}}, \tilde{A}_{\mathcal{C}}$ and $A_{\mathcal{C}}$, respectively.

Using the additive property of KL-divergence [16], from (25) and (26) we obtain that

$$\begin{aligned} D_{KL}(\text{View}_{\mathcal{C}}(A), \text{View}_{\mathcal{C}}(B)) &= D_{KL}(f_{\bar{A}_{\mathcal{H}}|A_{\mathcal{H}}}, f_{\bar{B}_{\mathcal{H}}|B_{\mathcal{H}}}) \\ &+ D_{KL}(f_{(\tilde{A}_{\mathcal{C}}, R_{\mathcal{C}})|A_{\mathcal{C}}}, f_{(\tilde{B}_{\mathcal{C}}, R_{\mathcal{C}})|B_{\mathcal{C}}}). \end{aligned} \quad (27)$$

As the affine coefficients $A_{\mathcal{C}}$ and $B_{\mathcal{C}}$ are identical to each other, we get from (22), the conditional probability distributions $f_{(\tilde{A}_{\mathcal{C}}, R_{\mathcal{C}})|A_{\mathcal{C}}}$ and $f_{(\tilde{B}_{\mathcal{C}}, R_{\mathcal{C}})|B_{\mathcal{C}}}$ are equivalent. Therefore,

$D_{KL}(f_{(\tilde{A}_{\mathcal{C}}, R_{\mathcal{C}})|A_{\mathcal{C}}}, f_{(\tilde{B}_{\mathcal{C}}, R_{\mathcal{C}})|B_{\mathcal{C}}}) = 0$. Upon substituting this in (27) we obtain that

$$D_{KL}(\text{View}_{\mathcal{C}}(A), \text{View}_{\mathcal{C}}(B)) = D_{KL}(f_{\bar{A}_{\mathcal{H}}|A_{\mathcal{H}}}, f_{\bar{B}_{\mathcal{H}}|B_{\mathcal{H}}}). \quad (28)$$

Let $\mathcal{G}_{\mathcal{H}} = (\mathcal{H}, \mathcal{E}_{\mathcal{H}})$ be the residual honest graph, and let $\mathcal{L}_{\mathcal{H}}$ denote the graph-Laplacian of $\mathcal{G}_{\mathcal{H}}$. As we assume that \mathcal{C} is not a vertex cut of \mathcal{G} , $\mathcal{G}_{\mathcal{H}}$ is connected. Therefore, substituting from Lemma 2 in (28) we obtain that

$$D_{KL}(\text{View}_{\mathcal{C}}(A), \text{View}_{\mathcal{C}}(B)) \leq \frac{1}{4\sigma^2 \underline{\mu}(\mathcal{L}_{\mathcal{H}})} \|A_{\mathcal{H}} - B_{\mathcal{H}}\|^2.$$

As $A_{\mathcal{C}} = B_{\mathcal{C}}$, we have $\|A_{\mathcal{H}} - B_{\mathcal{H}}\|^2 = \|A - B\|^2$ completing the proof. \blacksquare




Fig. 1. The agents' local errors from the optimizer (1) converges to zero.




Fig. 2. The p.d.f's of agent 1's affine coefficients generated numerically for two different scenarios; (i) the agents' private coefficients are A , and (ii) the agents' private coefficients are B .

V. NUMERICAL SIMULATION

In this section, we present a numerical simulation of the FS protocol. We consider a network of 3 agents, $\{1, 2, 3\}$, connected in a *complete* graph. The agents' private local costs are $h_1(x) = x^2 + x$, $h_2(x) = x^2 + 2x$, and $h_3(x) = x^2 + 3x$, where $x \in [-100, 100]$. Thus, $A = [\alpha_1, \alpha_2, \alpha_3] = [1, 2, 3]$. For computing the effective cost functions, defined in (9), the agents use $\sigma = 1$ in phase I. In phase II, we simulate the DGD on the effective cost functions. The absolute differences of the agents' local estimates from the minimizer of the aggregate cost is plotted in Fig. 1, for both the FS protocol and the conventional DGD algorithm, to show convergence.

We assume agent 3 to be corrupted by a passive adversary, i.e., $\mathcal{C} = \{3\}$ and $\mathcal{H} = \{1, 2\}$. We consider an alternate scenario where agents' affine coefficients are given by $B = [\beta_1, \beta_2, \beta_3] = [2, 1, 3]$. Note that $\alpha_3 = \beta_3$ and $\sum_{i=1}^3 \beta_i = \sum_{i=1}^3 \alpha_i = 6$. We simulate 100,000 executions of the FS protocol for both scenarios. The p.d.f's of agent 1's effective affine coefficients generated in phase I for both the scenarios are shown in Fig. 2. To compute the value of $D_{KL}(\text{View}_{\mathcal{C}}(A), \text{View}_{\mathcal{C}}(B))$, we first numerically approximate p_A and p_B , the respective conditional p.d.f.s of the effective coefficients $[\bar{\alpha}_1, \bar{\alpha}_2]$ and $[\bar{\beta}_1, \bar{\beta}_2]$ (defined by (23)) given the agents coefficients A and B , using the MATLAB's 'fitdist' function. Note that, owing to (28), $D_{KL}(\text{View}_{\mathcal{C}}(A), \text{View}_{\mathcal{C}}(B)) = D_{KL}(p_A, p_B)$. We obtain that p_A and p_B are Gaussian distributions with mean values $\mu_A = [1.00, 2.00]$ and $\mu_B = [2.00, 1.00]$, respectively, and an identical covariance matrix $\Sigma = [2.00, -2.00; -2.00, 2.00]$. Thus, $D_{KL}(p_A, p_B) = 0.5(\mu_A - \mu_B)\Sigma^\dagger(\mu_A - \mu_B)^T = 0.25$. This matches the theoretical bound computed by substituting $\mu(\mathcal{L}_{\mathcal{H}}) = 2$, $\sigma = 1$, and $\|A - B\|^2 = 2$ in Theorem 1.

VI. CONCLUDING REMARKS

We have presented a theoretical privacy analysis for the Function Sharing or FS protocol, a distributed optimization protocol proposed in [4] for protecting privacy of agents' costs against a passive adversary that corrupts some of the agents in the network. We have shown that the FS protocol preserves the statistical privacy of the polynomial terms of the honest agents' costs if the corrupted agents do not constitute a vertex cut of the network. Moreover, if the network has $(t + 1)$ -connectivity then the derived statistical privacy guarantee of the FS protocol holds true against *any* passive adversary that corrupts at most t agents.

REFERENCES

- [1] T. Yang, X. Yi, J. Wu, Y. Yuan, D. Wu, Z. Meng, Y. Hong, H. Wang, Z. Lin, and K. H. Johansson, "A survey of distributed optimization," *Annual Reviews in Control*, vol. 47, pp. 278 – 305, 2019.
- [2] S. Gade and N. H. Vaidya, "Private learning on networks," *arXiv preprint arXiv:1612.05236*, 2016.
- [3] M. C. Silaghi and D. Mitra, "Distributed constraint satisfaction and optimization with privacy enforcement," in *International Conference on Intelligent Agent Technology*. IEEE, 2004, pp. 531–535.
- [4] S. Gade and N. H. Vaidya, "Private optimization on networks," in *2018 American Control Conference (ACC)*. IEEE, 2018, pp. 1402–1409.
- [5] N. Gupta, J. Katz, and N. Chopra, "Privacy in distributed average consensus," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 9515–9520, 2017.
- [6] N. Gupta, "Privacy in distributed multi-agent collaboration: Consensus and optimization," Ph.D. dissertation, A. James Clark School of Engineering, University of Maryland - College Park, 2018.
- [7] E. Nozari, P. Tallapragada, and J. Cortés, "Differentially private distributed convex optimization via functional perturbation," *IEEE Control Netw. Syst.*, vol. 5, no. 1, pp. 395–408, 2018.
- [8] Z. Huang, S. Mitra, and N. Vaidya, "Differentially private distributed optimization," in *Proceedings of the 2015 International Conference on Distributed Computing and Networking*. ACM, 2015.
- [9] Y. Hong, J. Vaidya, N. Rizzo, and Q. Liu, "Privacy preserving linear programming," *arXiv preprint arXiv:1610.02339*, 2016.
- [10] Y. Lu and M. Zhu, "Privacy preserving distributed optimization using homomorphic encryption," *Automatica*, vol. 96, pp. 314 – 325, 2018.
- [11] C. Zhang, M. Ahmad, and Y. Wang, "Admm based privacy-preserving decentralized optimization," *IEEE Transactions on Information Forensics and Security*, vol. 14, no. 3, pp. 565–580, 2019.
- [12] J. Katz and Y. Lindell, *Introduction to modern cryptography*. CRC press, 2014.
- [13] N. Gupta, J. Katz, and N. Chopra, "Information-theoretic privacy in distributed average consensus," *arXiv:1809.01794*, 2018.
- [14] —, "Statistical privacy in distributed average consensus on bounded real inputs," in *2019 American Control Conference (ACC)*. IEEE, 2019, pp. 1836–1841.
- [15] S. Gade, J. Liu, and N. H. Vaidya, "A private and finite-time algorithm for solving a distributed system of linear equations," *arXiv preprint arXiv:2004.04680*, 2020.
- [16] T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, 2012.
- [17] S. Gade and N. H. Vaidya, "Distributed optimization of convex sum of non-convex functions," *arXiv preprint arXiv:1608.05401*, 2016.
- [18] A. Nedic and A. Ozdaglar, "Distributed subgradient methods for multi-agent optimization," *IEEE Transactions on Automatic Control*, vol. 54, no. 1, pp. 48–61, 2009.
- [19] C. Godsil and G. Royle, "Algebraic graph theory, volume 207 of graduate texts in mathematics," 2001.
- [20] S. Gade and N. H. Vaidya, "Private learning on networks: Part ii," *arXiv preprint arXiv:1703.09185*, 2017.
- [21] F. Yan, S. Sundaram, S. Vishwanathan, and Y. Qi, "Distributed autonomous online learning: Regrets and intrinsic privacy-preserving properties," *IEEE Transactions on Knowledge and Data Engineering*, vol. 25, no. 11, pp. 2483–2493, 2013.
- [22] I. Gutman and W. Xiao, "Generalized inverse of the laplacian matrix and some applications," *Bulletin de l'Academie Serbe des Sciences at des Arts (Cl. Math. Natur.)*, vol. 129, pp. 15–23, 2004.
- [23] C. R. Rao, *Linear statistical inference and its applications*. Wiley New York, 1973, vol. 2.