

# UC Irvine

## UC Irvine Previously Published Works

### Title

Cloud-Based Quadratic Optimization With Partially Homomorphic Encryption

### Permalink

<https://escholarship.org/uc/item/4cn9j7q5>

### Journal

IEEE Transactions on Automatic Control, 66(5)

### ISSN

0018-9286

### Authors

Alexandru, Andreea B  
Gatsis, Konstantinos  
Shoukry, Yasser  
[et al.](#)

### Publication Date

2021

### DOI

10.1109/tac.2020.3005920

Peer reviewed

# Cloud-based Quadratic Optimization with Partially Homomorphic Encryption

Andreea B. Alexandru, *Student Member, IEEE*, Konstantinos Gatsis, *Member, IEEE*,  
Yasser Shoukry, *Member, IEEE*, Sanjit A. Seshia, *Fellow, IEEE*, Paulo Tabuada, *Fellow, IEEE*,  
and George J. Pappas, *Fellow, IEEE*

**Abstract**—The development of large-scale distributed control systems has led to the outsourcing of costly computations to cloud-computing platforms, as well as to concerns about privacy of the collected sensitive data. This paper develops a cloud-based protocol for a quadratic optimization problem involving multiple parties, each holding information it seeks to maintain private. The protocol is based on the projected gradient ascent on the Lagrange dual problem and exploits partially homomorphic encryption and secure multi-party computation techniques. Using formal cryptographic definitions of indistinguishability, the protocol is shown to achieve computational privacy, i.e., there is no computationally efficient algorithm that any involved party can employ to obtain private information beyond what can be inferred from the party’s inputs and outputs only. In order to reduce the communication complexity of the proposed protocol, we introduced a variant that achieves this objective at the expense of weaker privacy guarantees. We discuss in detail the computational and communication complexity properties of both algorithms theoretically and also through implementations. We conclude the paper with a discussion on computational privacy and other notions of privacy such as the non-unique retrieval of the private information from the protocol outputs.

## I. INTRODUCTION

The recent push towards increasing connectivity of devices is enabling control applications to span wider geographical areas such as cities or even countries as in the case of smart grids. This increase in the number of available sensors, actuators and data has led to an increase in the controllers’ required computational capacity for such applications. One solution that has gained important traction is to outsource the computations to powerful remote servers, generically called clouds. Cloud-based computation services offer the potential to aggregate and to process information that is gathered from a large number of distributed agents, in domains such as machine learning, signal processing, smart grids, autonomous vehicles and control over the Internet-of-Things, with less overhead than distributed computations. Notwithstanding, there are substantial challenges that arise from the privacy-sensitive nature

of some of the collected data. Secure communication to and from the cloud is a crucial aspect, but is not sufficient to protect the users’ data. Recent examples such as data leakage and abuse by aggregating servers [1]–[3] have drawn attention to the risks of storing data in the clear and urged measures against an untrustworthy cloud. A plethora of different schemes addressing this challenge have been proposed in the literature, spanning many research areas.

Schemes that allow secure function evaluation on private inputs from multiple parties while preserving the privacy of the data involved have been a desirable technique for many decades, coining the term *secure multi-party computation*. Some secure multi-party computation protocols provide *computational security*, while other provide *information theoretic security*. Computational security guarantees that there is no computationally efficient algorithm the parties involved in the protocol can implement so that they gain private information beyond what can be inferred from their inputs and outputs of the protocol only. On the other hand, information theoretic security guarantees that no algorithms of unlimited computational power exist that can lead to private data leakage. One approach within this context is to encrypt the computation itself. This primarily refers to Yao’s *garbled circuits*, introduced in [4], and the protocol of Goldreich-Micali-Wigderson, introduced in [5], where both allow secure evaluation of functions represented as Boolean circuits. Although powerful with respect to the functions they can evaluate, garbled circuits require constructing high depth circuits and transmitting them over a communication channel, which incurs a substantial cost in computation and communication. While many improvements in the efficiency of garbled circuits have been made over the years, the complexity still prohibits the usage of such schemes in real time applications [6].

Another path within secure multi-party computation with computational security guarantees is computing on encrypted data. This approach, suggested in [7] under the name “privacy homomorphisms”, is currently known as homomorphic encryption. Homomorphic encryption allows evaluating specific types of functionals on encrypted inputs. Homomorphic encryption comes in three flavors: *fully*, *somewhat* and *partially* homomorphic encryption. Fully homomorphic encryption allows evaluation of Boolean functions over encrypted data and was introduced in [8], and further improved in terms of efficiency in [9]–[12]. However, the computational overhead is prohibitive, due to the complexity of the cryptosystem primitives and size of the encrypted messages [13]. Somewhat

A. B. Alexandru, K. Gatsis, G. J. Pappas are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104, USA, e-mail: aandreea,kgatsis,pappasg@seas.upenn.edu.

Y. Shoukry is with the Department of Electrical Engineering and Computer Science, University of California, Irvine, CA 92697, USA; e-mail: yshoukry@uci.edu.

S. A. Seshia is with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720, USA, e-mail: sseshia@eecs.berkeley.edu.

P. Tabuada is with the Department of Electrical Engineering, University of California, Los Angeles, Los Angeles, CA 90095, USA, e-mail: tabuada@ee.ucla.edu.

homomorphic encryption schemes, e.g. [14], are less complex, but support only a limited number of operations (additions and multiplications). Partially homomorphic encryption schemes are tractable but can support either multiplications between encrypted data, such as El Gamal [15], unpadded RSA [7] or additions between encrypted data, such as Paillier [16], Goldwasser-Micali [17], DGK [18].

While the above techniques ensure that the output is correctly computed without revealing anything about the inputs in the process, depending on the application, an adversary might be able to learn information about the input only from the output. Techniques such as *differential privacy* ensure that the revealed output does not break the privacy of the input. Differential privacy, introduced in [19] and surveyed in [20], [21], adds controlled randomness in the data, so that the results of statistical queries do not tamper with the privacy of the individual entries in a database, by introducing uncertainty in the output. However, the noise determines a trade-off between the accuracy of the results of the differentially private queries and the privacy parameter.

#### A. Related work

Homomorphic encryption has been recently used to design encrypted controllers, with partially homomorphic encryption schemes for linear controllers in [22], [23], model predictive control applications in [24], [25], with fully homomorphic encryption in [26], and to design observers in [27]. Unconstrained optimization problems are commonly used in machine learning applications, e.g., for linear regression, and several works have addressed gradient methods with partially homomorphic encryption [28], [29] or with ADMM in a distributed setting [30]. However, adding constraints substantially complicates the optimization problem by restricting the feasible domain. In addition, constraints introduce nonlinear operations in the optimization algorithm, which cannot be handled inherently by the partially homomorphic cryptosystems. In our prior work [31], [32], we showed how to overcome these difficulties by using blinded communication in a centralized setup. Recently, [33] proposed a distributed approach to the constrained optimization problem with homomorphic encryption, where each party locally performs the projection on the feasible domain on unencrypted data. Similarly, [34] proposed a distributed secure multi-party solution to regularized linear model training that also allows for malicious adversaries.

Examples of works that tackle private linear programming problems or machine learning applications with optimization subroutines using garbled circuits, secret sharing, and hybrid approaches between the above are [35]–[39] and the references within. In control applications, works such as [40]–[43] employ differential privacy to protect dynamic data. Furthermore, examples such as [44]–[46] make use of differential privacy techniques in optimization algorithms. Other lines of work solve optimization problems with non-cryptographic methods: using transformations, e.g. [47]–[50], which guarantee an uncertainty set around the private data, or using arbitrary functions for obfuscating the objective functions in a special graph topology [51].

#### B. Contributions

In this paper we consider a privacy-preserving cloud-based constrained quadratic optimization. In control theoretic applications and machine learning, linear and quadratic optimization problems arise frequently – e.g., state estimation under minimum square error, model predictive control, support vector machines – which require privacy guarantees of the data involved in the computation.

We develop a new tractable optimization protocol to privately solve constrained quadratic problems. Our protocol relies on cryptographic tools called encryption schemes. To solve the optimization problem on encrypted data, we use an additively homomorphic encryption scheme, where, in short, *addition commutes with the encryption function*. Thus, a party can process encrypted data without having direct access to the data. The novelty is how to handle in a privacy-preserving manner the constraints of the problem that introduce nonlinearities that cannot be supported by additively homomorphic encryption schemes. We show that a projected gradient method that operates on the Lagrange dual problem can alleviate this problem and can be efficiently run on encrypted data by exploiting communication between the participating parties.

The main contributions of our work are the following:

- We formally state and prove computational security guarantees for such protocol. The proof relies on applying cryptographic tools to the specific optimization algorithm that runs over multiple iterations.
- We propose an alternative protocol which sacrifices some privacy but involves less communication overhead.
- We implement the protocols and show the computational and communication complexity produce reasonable running times.

Furthermore, we emphasize the difference between the computational security guarantee with the *non-unique retrieval* guarantee that is important in such optimization problems. We finally point out that the present manuscript provides detailed security proofs and analyses not available in previous versions of the papers [31], [32].

*Organization.* The rest of the paper is organized as follows: we formulate the problem in Section II and formally introduce the computational security guarantees in Section III. We describe the cryptographic tools we use in Section IV, more specifically, the partially homomorphic scheme, along with its properties and the symmetric encryption used. In Section V, we describe the optimization theory behind the proposed protocol and justify the choice of the gradient method. Furthermore, we present the main protocol and the subroutines that compose it and in Section VI, we show that the protocol achieves privacy of the agent’s data and of the target node’s result. We discuss possible relaxations of the security guarantees and propose a more efficient protocol under these weaker conditions in Section VII. In Section VIII, we provide a privacy analysis of the problem concerning the input-output relation. We present the complexity analysis of the protocols and show that the experimental results point out reasonable running times in Section IX. Additional tools necessary for proving the security of the proposed protocol are introduced in Appendix A and further details on the protocols

used are given in Appendix B. Finally, the detailed privacy proofs are given in Appendix C.

*Notation.* We denote matrices by capital letters.  $\mathbb{R}^{m \times n}$  represents the set of  $m \times n$  matrices with real elements and  $\mathbb{S}_{++}^n$  represents the set of symmetric positive definite  $n \times n$  matrices. Element-wise inequalities between vectors are denoted by  $\preceq$ . The notation  $x \leftarrow X$  means that  $x$  is uniformly drawn from a distribution  $X$ . Pr denotes the probability taken over a specified distribution. We refer to algorithms that are run interactively by multiple parties as protocols.

## II. PROBLEM SETUP

### A. Motivating Examples

Quadratic optimization is a class of problems frequently employed in control system design and operation. As a first motivating example, consider the problem of estimating the state of a dynamical system from privacy-sensitive sensor measurements. For instance, such a problem arises in smart houses where the temperature or energy readings of the sensors are aggregated by a cloud controller and can reveal whether the owners are in the building. In particular, let the system dynamics and sensor measurements be described by:

$$\begin{aligned} x_{t+1} &= A x_t + w_t \\ y_t &= C x_t + v_t, \end{aligned} \quad (1)$$

for  $t = 0, \dots, T-1$ , where  $w_t, v_t$  are process and measurement noise respectively. The system and sensor parameters  $A \in \mathbb{R}^{n \times n}, C \in \mathbb{R}^{p \times n}$  can be thought as publicly available information while the sensor measurements  $y_0, \dots, y_{T-1}$  are privacy-sensitive data. The untrusted cloud has to collect the measurements and output an initial state estimate  $x_0$  of the system to a target agent, while maintaining the privacy of the sensor data and final output. A simple state estimate may be found as the solution to the least squares problem:

$$\min_{x_0 \in \mathbb{R}^n} \frac{1}{2} \sum_{t=0}^T \|y_t - C A^t x_0\|_2^2 = \frac{1}{2} \|y - \mathcal{O} x_0\|_2^2, \quad (2)$$

where

$$y = \begin{bmatrix} y_0 \\ y_1 \\ \vdots \\ y_{T-1} \end{bmatrix}, \quad \mathcal{O} = \begin{bmatrix} C \\ C A \\ \vdots \\ C A^{T-1} \end{bmatrix} \quad (3)$$

and  $\mathcal{O} \in \mathbb{R}^{Tp \times n}$  is the well-known system observability matrix. More general state estimation problems may also include constraints, e.g., the initial state may lie in a polyhedral set  $D x_0 \preceq b$  where the shape of the polyhedron captured by the matrix  $D$  is publicly known but its position and size captured by the vector  $b$  is private information.

As a second example, consider a control problem with privacy-sensitive objectives. Suppose we are interested in steering a dynamical system:

$$\begin{aligned} x_{t+1} &= A x_t + B u_t \\ y_t &= C x_t, \end{aligned} \quad (4)$$

starting from a private initial position  $x_0$  while guaranteeing safety state constraints of the form:

$$\underline{x}_t \preceq x_t \preceq \bar{x}_t, t = 1, \dots, T, \quad (5)$$

where  $\underline{x}_t, \bar{x}_t$  for  $t = 1, \dots, T$  are private values. Such problems arise when exploring vehicles are deployed in uncertain or hazardous environments or when different users compete against each other and want to hide their tactics.

Denote by  $x^r$  and  $u^r$  the private values such that the system is guaranteed to track a reference  $r$ . Then, we need to solve the problem:

$$\begin{aligned} \min_{x_k \in \mathbb{R}^n, u_k \in \mathbb{R}^m} & \frac{1}{2} \sum_{t=0}^{T-1} \|y_t - C x_t^r\|_Q^2 + \|u_t - u_t^r\|_R^2 \\ \text{s.t.} & x_{t+1} = A x_t + B u_t, \quad y_t = C x_t \\ & \underline{x}_t \preceq x_t \preceq \bar{x}_t, \quad t = 1, \dots, T, \end{aligned} \quad (6)$$

where  $Q, R$  are positive definite cost matrices and we used the quadratic norm notation  $\|z\|_P^2 := z^\top P z$ .

The publicly known system dynamics (4) are equality constraints which can be eliminated to obtain a control input design problem with only linear inequality constraints:

$$\begin{aligned} \min_{x_k \in \mathbb{R}^n, u_k \in \mathbb{R}^m} & \frac{1}{2} \sum_{t=0}^{T-1} u_t^\top H_t u_t + [x_0 \quad x_t^r \quad u_t^r] F_t u_t \\ \text{s.t.} & \underline{x}_t \preceq A^t x_0 + \sum_{j=0}^{t-1} A^{t-1-j} B u_j \preceq \bar{x}_t, \\ & t = 1, \dots, T, \end{aligned} \quad (7)$$

with  $H_t$  and  $F_t$  are matrices that depend on the costs  $Q, R$  and the system's dynamics (4) appropriately computed from (6).

### B. Problem statement

The above examples can be modeled as constrained quadratic optimization problems with distributed private data. We consider three types of parties involved in the problem: a number of agents  $\mathcal{A}_i, i = 1, \dots, p$ , a cloud server  $\mathcal{C}$  and a target node  $\mathcal{T}$ . The purpose of this setup is to solve an optimization problem with the data provided by the agents and the computation performed on the cloud and send the result to the target node. The architecture is presented in Figure 1.

Let us consider a strictly-convex quadratic optimization problem, which we assume to be feasible:

$$\begin{aligned} x^* &= \arg \min_{x \in \mathbb{R}^n} \frac{1}{2} x^\top Q_C x + c_A^\top x \\ \text{s.t.} & A_C x \preceq b_A, \end{aligned} \quad (8)$$

where the variables and the parties to which they belong to are described as follows:

**Agents**  $\mathcal{A} = (\mathcal{A}_1, \dots, \mathcal{A}_p)$ : The agents are parties with low computational capabilities that possess the private information  $b_A$  and  $c_A$ . The private information is decomposed across the agents as:  $b_A = (b_1, \dots, b_p)$  and  $c_A = (c_1, \dots, c_p)$ , with  $b_i \in \mathbb{R}^{m_i}$  and  $c_i \in \mathbb{R}^{n_i}$  being the private data of agent  $i$  such that  $\sum_{i=1}^p m_i = m$  and  $\sum_{i=1}^p n_i = n$ .

**Cloud**  $\mathcal{C}$ : The cloud is a party with high computational capabilities that has access to the matrices  $Q_C \in \mathbb{S}_{++}^n$  and

$A_C \in \mathbb{R}^{m \times n}$ . When the computation is sophisticated and/or involves proprietary algorithms,  $Q_C$  and  $A_C$  are private data of the cloud. In order to capture a greater number of problems, we will also consider the case where  $Q_C$  or  $A_C$  are public.

**Target Node  $\mathcal{T}$ :** The target node is a party with more computational capabilities than the agents that is interested in the optimal solution  $x^*$  of the problem. The target node can be either one of the agents or a separate cloud server.




Fig. 1. Architecture of the problem: Agents are low-resource parties that have private data that they outsource to a powerful server, called the cloud. The cloud has to solve an optimization problem on the private data of the agents and send the result to a party called the target node, which will help with the computations.

Note that in the first motivating example (2), the matrix  $Q_C$  in (8) corresponds to the publicly known matrix  $\mathcal{O}^T \mathcal{O}$  and the vector  $c_A$  corresponds to the private vector  $-\mathcal{O}^T y$ . For the objective to be strongly convex, we require that  $\text{rank}(\mathcal{O}) = n$ , which also implies standard system observability. Similarly, in the second example (7), the matrix  $Q_C$  is composed from the regulating cost matrices, the cost vector  $c_A$  is formed from the private initial conditions and steady-state solution for the reference tracking, mixed by the system's dynamics, and the constraints vector  $b_A$  depends on the private state bounds and initial condition.

### C. Adversarial model

In most cloud applications, the service provider has to deliver the contracted service or otherwise the clients would switch to another service provider. This means that the cloud cannot alter the data it receives. Moreover, the clients' interest is to obtain the correct result from the service they pay for, hence, the agents and target node will also not alter the data. However, the parties can locally process the data they receive in any fashion they want. This model is known as semi-honest, which is defined as follows:

**Definition 1. (Semi-honest model)** A party is semi-honest if it does not deviate from the steps of the protocol but may store the transcript of the messages exchanged and process the data received in order to learn more information than stipulated by the protocol.

This model also holds when considering eavesdroppers on the communication channels. An adversarial model that only considers eavesdroppers as adversaries differs from the semi-honest model by the fact that, apart from the untrusted channels, the parties that perform the computations are also

not trusted. Malicious and active adversaries – that diverge from the protocols or tamper with the messages – are not considered in this paper.

The purpose of the paper is to solve Problem (8) using a secure multi-party computation protocol for semi-honest parties. This protocol takes as inputs the private data of the agents, as well as the cloud's data, and involves the parties in exchanging messages and participating in some specified computations, and eventually outputs to the target node the solution of the optimization problem. This protocol should guarantee *computational security*. More specifically, the cloud cannot obtain any information about the private inputs of the agents and the output of the target node and similarly, the target node cannot obtain any information about the private inputs of the agents and the cloud, other than what they can compute using their inputs and outputs and public information, by running a computationally efficient algorithm after the execution of the protocol.

### III. PRIVACY GOALS

We introduce some notion necessary to the security definitions. In what follows,  $\{0, 1\}^*$  defines a sequence of bits of unspecified length. An ensemble  $X = \{X_n\}_{n \in \mathbb{N}}$  is a sequence of random variables ranging over strings of bits of length polynomial in  $n$ , denoted by  $\text{poly}(n)$ , arising from distributions defined over a finite set  $\Omega$ . By positive polynomials, we refer to the set of polynomials that on positive inputs have positive values.

**Definition 2. (Statistical indistinguishability [52, Ch. 3])** The ensembles  $X = \{X_n\}_{n \in \mathbb{N}}$  and  $Y = \{Y_n\}_{n \in \mathbb{N}}$  are **statistically indistinguishable**, denoted  $\stackrel{s}{\equiv}$ , if for every positive polynomial  $p$ , and all sufficiently large  $n$ :

$$\frac{1}{2} \sum_{\alpha \in \Omega} |\Pr[X_n = \alpha] - \Pr[Y_n = \alpha]| < \frac{1}{p(n)},$$

where the quantity on the left is called the statistical distance between the two ensembles.

It can be proved that two ensembles are statistically indistinguishable if no algorithm can distinguish between them. Computational indistinguishability is a weaker notion of the statistical version, as follows:

**Definition 3. (Computational Indistinguishability [52, Ch. 3])** The ensembles  $X = \{X_n\}_{n \in \mathbb{N}}$  and  $Y = \{Y_n\}_{n \in \mathbb{N}}$  are **computationally indistinguishable**, denoted  $\stackrel{c}{\equiv}$ , if for every **probabilistic polynomial-time** algorithm  $D : \{0, 1\}^* \rightarrow \{0, 1\}$ , called the distinguisher, every positive polynomial  $p$ , and all sufficiently large  $n$ , the following holds:

$$|\Pr_{x \leftarrow X_n}[D(x) = 1] - \Pr_{y \leftarrow Y_n}[D(y) = 1]| < \frac{1}{p(n)}.$$

The previous definition is relevant when defining the privacy goals to be guaranteed by the protocol: **two-party privacy** of the sensitive data of the parties. The intuition is that a protocol privately computes a functionality if nothing is learned after its execution, i.e., if all information obtained by a party after the execution of the protocol (while also keeping a record of

the intermediate computations) can be obtained only from the inputs and outputs available to that party.

**Definition 4.** (Two-party privacy w.r.t. semi-honest behavior [53, Ch. 7]) Let  $f : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}^* \times \{0, 1\}^*$  be a functionality, and  $f_1(x_1, x_2), f_2(x_1, x_2)$  denote the first and second components of  $f(x_1, x_2)$ , for any inputs  $x_1, x_2 \in \{0, 1\}^*$ . Let  $\Pi$  be a two-party protocol for computing  $f$ . The **view** of the  $i$ -th party ( $i = 1, 2$ ) during an execution of  $\Pi$  on the inputs  $(x_1, x_2)$ , denoted  $V_i^\Pi(x_1, x_2)$ , is  $(x_i, \text{coins}, m_1, \dots, m_t)$ , where *coins* represents the outcome of the  $i$ 'th party's internal coin tosses, and  $m_j$  represents the  $j$ -th message it has received. For a deterministic functionality  $f$ , we say that  $\Pi$  **privately computes**  $f$  if there exist probabilistic polynomial-time algorithms, called **simulators**, denoted  $S_i$ , such that:

$$\{S_i(x_i, f_i(x_1, x_2))\}_{x_{1,2} \in \{0,1\}^*} \stackrel{c}{\equiv} \{V_i^\Pi(x_1, x_2)\}_{x_{1,2} \in \{0,1\}^*}.$$

This definition assumes the correctness of the protocol, i.e., the probability that the output of the parties is not equal to the result of the functionality applied to the inputs is negligible [54].

For protocols that involve more than two parties, Definition 4 can be extended to **multi-party privacy**, by also taking into consideration coalitions of semi-honest parties. The next definition states that multi-party protocol privately computes the functionality it runs if everything that a coalition of parties can obtain from participating in the protocol and keeping records of the intermediate computations can be obtained only from the inputs and outputs of these parties.

**Definition 5.** (Multi-party privacy w.r.t. semi-honest behavior [53, Ch. 7]) Let  $f : (\{0, 1\}^*)^n \rightarrow (\{0, 1\}^*)^n$  be a  $n$ -ary functionality, where  $f_i(x_1, \dots, x_n)$  denotes the  $i$ -th element of  $f(x_1, \dots, x_n)$ . Denote the inputs by  $\bar{x} = (x_1, \dots, x_n)$ . For  $I = \{i_1, \dots, i_t\} \subset [n] = \{1, \dots, n\}$ , we let  $f_I(\bar{x})$  denote the subsequence  $f_{i_1}(\bar{x}), \dots, f_{i_t}(\bar{x})$ , which models a coalition of a number of parties. Let  $\Pi$  be a  $n$ -party protocol that computes  $f$ . The **view** of the  $i$ -th party during an execution of  $\Pi$  on the inputs  $\bar{x}$ , denoted  $V_i^\Pi(\bar{x})$ , is defined as in Definition 4, and we let the view of a coalition be denoted by  $V_I^\Pi(\bar{x}) = (I, V_{i_1}^\Pi(\bar{x}), \dots, V_{i_t}^\Pi(\bar{x}))$ . For a deterministic functionality  $f$ , we say that  $\Pi$  **privately computes**  $f$  if there exist simulators  $S$ , such that, for every  $I \subset [n]$ , it holds that, for  $\bar{x}_t = (x_{i_1}, \dots, x_{i_t})$ :

$$\{S(I, (\bar{x}_t), f_I(\bar{x}_t))\}_{\bar{x} \in (\{0,1\}^*)^n} \stackrel{c}{\equiv} \{V_I^\Pi(\bar{x})\}_{\bar{x} \in (\{0,1\}^*)^n}.$$

Further underlying details related to the cryptographic notion of security and its interpretation are presented in Appendix A-A.

## IV. ENCRYPTION SCHEME

### A. Additively Homomorphic Encryption

Let  $E(\cdot)$  define a generic encryption primitive, with domain the space of private data, called plaintexts, and codomain the space of encrypted data, called ciphertexts. For probabilistic encryption schemes, the encryption primitive also takes as input a random number. The decryption primitive  $D(\cdot)$  is

defined on the space of ciphertexts and takes values in the space of plaintexts. Additively homomorphic schemes satisfy the property that there exists an operator  $\oplus$  defined on the space of ciphertexts such that:

$$E(a) \oplus E(b) \subset E(a + b), \quad (9)$$

for any plaintexts  $a, b$  supported by the scheme. Here, we use set inclusion instead of equality because the encryption of a message is not unique in probabilistic cryptosystems. Intuitively, equation (9) means that performing this operation on the two encrypted messages, we obtain a ciphertext that is equivalent to the encryption of the sum of the two plaintexts. Formally, the decryption primitive  $D(\cdot)$  is a homomorphism between the group of ciphertexts with the operator  $\oplus$  and the group of plaintexts with addition  $+$ , which justifies the name of the scheme. It is immediate to see that if a scheme supports addition between encrypted messages, it will also support subtraction, by adding the additive inverse, and multiplication between an integer plaintext and an encrypted message, obtained by adding the encrypted messages for the corresponding number of times.

In this paper, we use the popular Paillier encryption, but any other additively homomorphic encryption scheme can be employed. The **Paillier cryptosystem** [16] is an asymmetric additively homomorphic encryption scheme. Asymmetric or public key cryptosystems involve a pair of keys: a public key that is disseminated publicly, and which is used for the encryption of the private messages, and a private key which is known only to its owner, used for the decryption of the encrypted messages. In this section, we will only present the notation for the encryption and for the allowed operations between the encrypted messages and postpone the implementation details to Appendix A-B.

In the Paillier encryption scheme, the private data are elements of the ring of integers modulo  $N$ , denoted by  $\mathbb{Z}_N$ , where  $N$  is a large integer of  $\sigma$  bits, called the Paillier modulus. The encrypted messages take values in the positive integers modulo  $N^2$ , denoted by  $\mathbb{Z}_{N^2}^*$ . Paillier is a probabilistic encryption scheme, which means that the encryption primitive also takes as an argument a random number which, for simplicity, we avoid in this notation. For a plaintext message  $a$ , we denote the Paillier encryption by  $[[a]] := E(a)$ , for an instance of the random number. For readability, we will use throughout the paper the following abstract notation for the operations on the encrypted space:

$$\begin{aligned} [[a]] \oplus [[b]] &\stackrel{d}{=} [[a + b]] \\ b \otimes [[a]] &\stackrel{d}{=} [[ba]], \quad \text{for plaintexts } a, b \in \mathbb{Z}_N, \end{aligned} \quad (10)$$

where  $\stackrel{d}{=}$  means that the equality holds after applying the decryption primitive on both sides [16]. The expressions for the abstracted operations are described in Appendix A-B. We will slightly abuse these notations to denote encryptions of vectors and additions and multiplication by vectors and matrices. In our protocol, the target node  $\mathcal{T}$  is the owner of a pair of Paillier keys. Hence, for everything that follows,  $[[\cdot]]$  denotes the encryption with the target node's public key  $pk_{\mathcal{T}}$ .

*Example: Private unconstrained quadratic optimization.* In order to illustrate the usage of the Paillier encryption scheme for computations on encrypted data, consider the following example of an unconstrained optimization problem:

$$\min_x f(x) = \frac{1}{2}x^\top Q_C x + c_A^\top x, \quad (11)$$

where the variables have the same meaning as in Problem (8) described in Section II. The cloud has access to the matrix  $Q_C$ . If the cloud also had access to private data  $c_A$  from the agents, the gradient method could be employed to obtain the optimal unencrypted solution of Problem (11). An iteration of the gradient method has the form:

$$x_{k+1} = x_k - \eta \nabla f(x_k) = x_k - \eta Q_C x_k - \eta c_A,$$

where  $\eta > 0$  is the step-size chosen by the cloud and  $k = 0, \dots, K-1$  for a number of iterations  $K$ . However, in order to protect their private data and the privacy of the solution  $x^*$ , the agents send the value  $c_A$  encrypted to the cloud as  $[[c_A]]$ . It is very important to notice that *for a quadratic objective function, only linear operations in the private data  $x_k$  and  $c_A$  are required in order to compute one iteration of the gradient descent algorithm.* Hence, by taking advantage of the additively homomorphic property (10), the cloud can locally compute the optimal solution by the gradient descent in the encrypted domain, for all  $k = 0, \dots, K-1$ :

$$[[x_{k+1}]] = (I - \eta Q_C) \otimes [[x_k]] \oplus (-\eta) \otimes [[c_A]], \quad (12)$$

and send the final point  $[[x_K]]$  to the target node to decrypt and obtain the desired result. Such a protocol satisfies the desired security according to Definition 4 provided in Section III by the fact that the cloud only has access to data encrypted by a semantically secure encryption scheme, such as Paillier [16]. More details are provided in Section VI.

However, the optimization problem considered in this paper is a quadratic problem with constraints, which introduce nonlinear operations in the optimization algorithm and cannot be handled inherently by the Paillier cryptosystem. We will describe in Section V how we leverage communication between the parties in order to achieve privacy-preserving nonlinear operations.

### B. Symmetric encryption scheme

Apart from the (public key) partially homomorphic cryptosystem above, we will also use a specific symmetric key cryptosystem for reasons that will become apparent in the protocol, related to hiding the data from the target node. This symmetric key cryptosystem is similar to the well-known one-time pad [55], [56], and can be thought of as additively blinding a private value by noise. In contrast to public key cryptography, symmetric key algorithms perform the encryption and decryption with the same key. For messages of  $l$  bits, the key is generated as  $sk \in \mathbb{Z}$  with length of  $\lambda + l$  bits, where  $\lambda$  is the security parameter. The encryption primitive is  $E'(a) = a + sk$ , with  $a \in [0, 2^l) \cap \mathbb{Z}$ , and the decryption is obtained as  $a = D'(E'(a)) = E'(a) - sk$ . This is also called two-out-of-two additive secret sharing [57] when the addition is on an abelian group. The security of this

scheme lies on generating a uniformly random key and on using this key for encryption only once, which yields that the distribution of  $E'(a)$  is statistically indistinguishable from a random number sampled of  $l + \lambda + 1$  bits. For this reason,  $\lambda$  is called the statistical security parameter. We also notice that *this symmetric cryptosystem is compatible with the Paillier cryptosystem, in the sense that the two encryptions commute:  $D'(E(E'(a))) = E(a)$  by using  $E(sk)$  instead of  $sk$  for decryption, where  $E'(a)$  is performed on the message space of the Paillier cryptosystem.* Such a scheme is commonly employed for blinding messages, for instance in [58]–[60].

For simplicity, in the protocols described in Section V, we will directly use the summation with a random number to depict the above symmetric key cryptosystem, respectively the difference by the same random number for decryption.

The strength of a cryptosystem relies on the computational intractability of retrieving the private key from the public information – an adversary holding the public information cannot find the private key by brute force computations. Public key cryptosystems like the Paillier cryptosystem involve a security parameter  $\sigma$ , which is usually the length of the modulus of the public key, and symmetric key cryptosystems involve a security parameter  $\lambda$ , which is the bit-length of the key. As is common in practice, increasing the size of the security parameters to values larger than, say 1024 bits for  $\sigma$  and 80 bits for  $\lambda$ , increases the security of the system against any computationally efficient inference attack.

## V. SECURE CONSTRAINED QUADRATIC OPTIMIZATION

In this section we introduce our main algorithmic results. Let us first describe the optimization algorithm used for solving the minimization problem (8) on unencrypted data.

For strongly convex problems, one can resort to duality theory [61, Ch. 5] to compute the projection on the feasible set, and be able to retrieve the optimal value of the primal problem from the optimal value of the dual problem. For the quadratic optimization problem (8), its dual is also a quadratic optimization problem:

$$\begin{aligned} \mu^* = \arg \max_{\mu \in \mathbb{R}^m} & -\frac{1}{2}(A_C^\top \mu + c_A)^\top Q_C^{-1}(A_C^\top \mu + c_A) - \mu^\top b_A \\ \text{s.t. } & \mu \succeq 0. \end{aligned} \quad (13)$$

The dual objective function is denoted by  $g(\mu)$  and its gradient is equal to:

$$\nabla g(\mu) = -A_C Q_C^{-1}(A_C^\top \mu + c_A) - b_A. \quad (14)$$

Under standard constraint qualifications, e.g., Slater's condition [61, Ch. 5], strong duality between the primal and dual holds, which means the optimal objective in the primal problem (8) is equal to the objective in the dual problem (13). Moreover, the optimality conditions (Karush-Kuhn-Tucker) hold and are the following:

$$Q_C x^* + A_C^\top \mu^* + c_A = 0 \quad (15)$$

$$A_C x^* - b_A \preceq 0, \quad \mu^* \succeq 0 \quad (16)$$

$$\mu_i^*(a_i^\top x^* - b_i) = 0, \quad i = 1, \dots, m. \quad (17)$$

For strictly convex problems, i.e.,  $Q_C \in \mathbb{S}_{++}^n$ , the optimal solution of the primal problem can be obtained from (15) as

$$x^* = -Q_C^{-1}(A_C^\top \mu^* + c_A).$$

An algorithm for computing the optimum in problem (13), which we will show is also compatible with the partially homomorphic encryption of Section IV, is the projected gradient ascent method. The projected gradient ascent is composed by iterations of the following type:

$$\mu_{k+1} = \max\{\mathbf{0}, \mu_k + \eta \nabla g(\mu_k)\}, \quad (18)$$

where  $\eta > 0$  is the step size and  $\mu_{k+1}$  is the projected value of  $\mu_k + \eta \nabla g(\mu_k)$  over the non-negative orthant. For full rank of  $A_C Q_C^{-1} A_C^\top$ , the dual problem is strictly convex and the algorithm converges with a linear rate [62] for a fixed step size  $\eta = \frac{1}{L}$ , where  $L = \lambda_{\max}(A_C Q_C^{-1} A_C^\top)$ . For non-strictly convex dual function, the gradient ascent algorithm converges in sublinear time [62].

#### A. Projected gradient ascent on encrypted data

As stated in Section II, we aim to solve an optimization problem outsourced to the cloud on private distributed data from the agents and send the result to the target node. To protect the agents' data, we use an encryption scheme that allows the cloud to perform linear manipulations on encrypted data, as described in Section IV. To this end, the target node generates a pair of keys  $(pk_{\mathcal{T}}, sk_{\mathcal{T}})$  and distributes the public key to the agents and the cloud, enabling them to encrypt their data, which only the target node will be able to decrypt, using the private key. We consider that all the data is represented on integers of  $l$  bits and comment on this further in Section V-E.

The main difficulty in performing the projected gradient ascent on encrypted data is performing iteration (18). We have already seen in the example in Section IV-A that the update of the iterate in the direction of the gradient ascent can be computed locally by the cloud directly on the encrypted data (12). However a first challenge lies in performing the comparison with zero. Due to the probabilistic nature of the Paillier encryption scheme, *the order on the plaintext space is not preserved on the ciphertext space* and comparison on encrypted data cannot be performed locally by the cloud. Moreover, after the comparison is performed, the update of the encrypted iterate (18) has to be done in a private way, so that the result of the maximum operation is not revealed to any of the parties involved (the cloud and the target node). These two steps are the main computational bottleneck in the protocol we propose, as both require secure communication between the cloud and the target node. A preliminary version of this solution was presented in [32].

We can privately achieve the two steps mentioned above in three stages. First, the cloud has to randomize the order of the two encrypted variables it wants to compare (Protocol 1). Second, the cloud and target engage in an interactive comparison protocol that takes as inputs the two randomized variables and outputs the result of the comparison to the target node (Protocol 2). Third, the update of the dual iterate is achieved through an interactive protocol between the cloud

and target node, which takes as inputs the two randomized variables and the result of the comparison and outputs to the cloud the updated iterate (Protocol 3). Throughout this paper, by comparison we mean element-wise comparison, since the variable  $\mu$  is a vector.

#### B. Secure comparison protocol

In order to privately compute (18), i.e., hide the result from *all* the parties involved, we want to keep the result of the comparison of the updated iterate  $\mu_k + \eta \nabla g(\mu_k)$  with zero unknown to both the cloud and the target node. The comparison protocol will reveal the result of the comparison between the two inputs to the target node  $\mathcal{T}$ . However, if we introduce an additional step where  $\mathcal{C}$  randomizes the order of the two values that it wants to compare, then  $\mathcal{T}$  does not learn any information by knowing the result of the comparison.

---

##### PROTOCOL 1: Randomization step

**Input:**  $\mathcal{C}$ :  $[[\bar{\mu}]], [[0]]$ , where  $\bar{\mu} := \mu + \eta \nabla g(\mu)$

**Output:**  $\mathcal{C}$ :  $[[a]], [[b]]$

- 1:  $\mathcal{C}$ : choose a random permutation  $\pi$  on two elements
  - 2:  $\mathcal{C}$ : output  $[[a]], [[b]] \leftarrow \pi([[0]], [[\bar{\mu}]])$
- 

Next, we will demonstrate how the comparison protocol works via the following instantiation. Damgård, Geisler and Krøigaard introduced a protocol in [18], [63] for secure comparison of two private inputs of different parties, which we will call the DGK protocol. To this end, they also propose an additively homomorphic encryption scheme, which has the property that checking if the value zero is encrypted is more efficient than simply decrypting, which is useful for comparisons and working with bits. An extension of this protocol to the case where none of the parties knows the two values to be compared, which is of interest to us, and some improvements in terms of efficiency were proposed in [64].

---

##### PROTOCOL 2: Protocol for secure two-party comparison with two encrypted inputs using DGK [18], [64]

**Input:**  $\mathcal{C}$ :  $[[a]], [[b]]$ ;  $\mathcal{T}$ :  $sk_{\mathcal{T}}, sk_{DGK}$

**Output:**  $\mathcal{T}$ : bit  $t$ :  $(t = 1) \Leftrightarrow (a \leq b)$

- 1:  $\mathcal{C}$ : choose random number  $\rho$
  - 2:  $\mathcal{C}$ :  $[[z]] \leftarrow [[b]] \oplus (-1) \otimes [[a]] \oplus [[2^l + \rho]]$ , send  $[[z]]$  to  $\mathcal{T}$   
 $\triangleright z \leftarrow b - a + 2^l + \rho$
  - 3:  $\mathcal{T}$ : decrypt  $[[z]]$
  - 4:  $\mathcal{C}$ :  $\alpha \leftarrow \rho \bmod 2^l$
  - 5:  $\mathcal{T}$ :  $\beta \leftarrow z \bmod 2^l$
  - 6:  $\mathcal{C}, \mathcal{T}$ : perform a comparison protocol, e.g, DGK, such that  $\mathcal{C}$  gets  $\delta_{\mathcal{C}}$  and  $\mathcal{T}$  gets  $\delta_{\mathcal{T}}$  with  $\delta_{\mathcal{C}} \vee \delta_{\mathcal{T}} = (\alpha \leq \beta)$   $\triangleright \vee$  denotes the exclusive or operation
  - 7:  $\mathcal{T}$ : encrypt  $[[z \div 2^l]]$  and  $[[\delta_{\mathcal{T}}]]$  and send them to  $\mathcal{C}$
  - 8:  $\mathcal{C}$ :  $[[t']] \leftarrow \begin{cases} [[\delta_{\mathcal{T}}]] & \text{if } \delta_{\mathcal{C}} = 1 \\ [[1]] \oplus (-1) \otimes [[\delta_{\mathcal{T}}]] & \text{otherwise} \end{cases}$
  - 9:  $\mathcal{C}$ :  $[[t]] \leftarrow [[z \div 2^l]] \oplus (-1) \otimes ([[\rho \div 2^l]] \oplus [[t']])$   
 $\triangleright t \leftarrow z \div 2^l - \rho \div 2^l - t'$
  - 10:  $\mathcal{C}$ : send  $[[t]]$  to  $\mathcal{T}$
  - 11:  $\mathcal{T}$ : decrypts  $[[t]]$
- 

The comparison protocol that will be used in our optimization protocol is as follows. Let  $\mathcal{C}$  have two encrypted values under the Paillier scheme  $[[a]]$  and  $[[b]]$  that it obtained



after running Protocol 1, and let  $\mathcal{T}$  have the decryption key. Furthermore, let  $\mathcal{T}$  also have the decryption key of the DGK homomorphic encryption scheme, which we describe in Appendix B. At the end of the protocol,  $\mathcal{T}$  will have the result of the comparison in the form of one bit  $t$  such that  $(t = 1) \Leftrightarrow (a \leq b)$ . Let  $l$  denote the number of bits of the unencrypted inputs  $a, b$ . Protocol 2 is based on the fact that the most significant bit of  $(b - a + 2^l)$  is the bit that indicates if  $(a \leq b)$ .

The security of this protocol is proved in [18], [63], [64] and the proof is omitted here.

**Proposition 1.** *Protocol 2 is secure in the semi-honest model, according to Definition 4.*

### C. Secure update protocol

Moreover, we need to ensure that when the cloud  $\mathcal{C}$  updates the value of the dual iterate at iteration  $k+1$  in equation (18), it does not know the new value. The solution is to make the cloud blind the values of  $[[a]]$  and  $[[b]]$  and send them to the target node in this order, where the latter selects the value accordingly to the comparison result and then sends it back to the cloud. However, there are two important issues that have to be addressed in order for the update step to not leak information about the sign of the iterate: the blinding should be additive and effectuated with different random values, and the ciphertexts should be refreshed. The reasons are the following: if the blinding is multiplicative, by decrypting the product, the target knows which one of the values is zero. Moreover, if the two values are additively blinded with the same random value, the target can subtract them and reveal at least if the value is zero. Re-randomization of the encryptions is necessary so that the cloud cannot simply compare  $[[a]]$  and  $[[b]]$  with the received value. This can be done by adding an encryption of zero or by decryption followed by encryption. Protocol 3 is the solution to the update problem:

---

#### PROTOCOL 3: Secure update of the dual variable

**Input:**  $\mathcal{C}$ :  $[[a]], [[b]]$ ;  $\mathcal{T}$ :  $t$  such that  $(t = 1) \Leftrightarrow (a \leq b)$

**Output:**  $\mathcal{C}$ :  $[[\mu]]$

- 1:  $\mathcal{C}$ : choose two random numbers  $r, s$
  - 2:  $\mathcal{C}$ :  $[[\bar{a}]] \leftarrow [[a]] \oplus [[r]], [[\bar{b}]] \leftarrow [[b]] \oplus [[s]]$
  - 3:  $\mathcal{C}$ : send  $[[\bar{a}]]$  and  $[[\bar{b}]]$  to  $\mathcal{T}$
  - 4: **if**  $t = 0$  **then**  $\mathcal{T}$ :  $[[v]] \leftarrow [[\bar{a}]] = [[\bar{a}]] + [[0]]$
  - 5: **else**  $\mathcal{T}$ :  $[[v]] \leftarrow [[\bar{b}]] = [[\bar{b}]] + [[0]]$
  - 6: **end if** ▷ Refresh the ciphertext
  - 7:  $\mathcal{T}$ : send  $[[v]]$  and  $[[t]]$  to  $\mathcal{C}$
  - 8:  $\mathcal{C}$ :  $[[\mu]] \leftarrow [[v]] \oplus r \otimes [[t]] \oplus [[-r]] \oplus (-s) \otimes [[t]]$   
▷  $\mu \leftarrow v + r(t - 1) - st$
- 

The intuition behind the protocol is that if  $a \leq b$ , then  $t = 1$  and we obtain  $\mu = \bar{b} - s = b$ , and otherwise,  $t = 0$  and we obtain  $\mu = \bar{a} - r = a$ . In both cases the cloud correctly updates the dual variable with the projected value.

### D. Protocol for solving strictly-convex quadratic problems

Having defined these protocols, we can now build a protocol that represents one iteration (18) of the dual projected gradient ascent method.

---

#### PROTOCOL 4: Secure iteration of the dual projected gradient ascent method

**Input:**  $\mathcal{C}$ :  $A_C \in \mathbb{R}^{m \times n}, Q_C \in \mathbb{S}_{++}^n, [[b_A]], [[c_A]], \eta > 0, [[\mu_k]]$ ;  $\mathcal{T}$ :  $sk_{\mathcal{T}}$

**Output:**  $\mathcal{C}$ :  $[[\mu_{k+1}]]$

- 1:  $\mathcal{C}$ :  $[[\nabla g(\mu_k)]] \leftarrow (-A_C Q_C^{-1} A_C^T) \otimes [[\mu_k]] \oplus (-A_C Q_C^{-1}) \otimes [[c_A]] \oplus (-1) \otimes [[b_A]]$   
▷ Compute the encrypted gradient as in (14)
  - 2:  $\mathcal{C}$ :  $[[\bar{\mu}_k]] \leftarrow [[\mu_k]] \oplus \eta \otimes [[\nabla g(\mu_k)]]$  ▷ Update the value in the ascent direction
  - 3:  $\mathcal{C}, \mathcal{T}$  truncate  $[[\bar{\mu}_k]]$  to  $l$  bits
  - 4:  $\mathcal{C}$  execute Protocol 1:  $\mathcal{C}$  gets  $[[a_k]], [[b_k]]$  ▷ Randomly assign  $[[a_k]], [[b_k]]$  with values of  $[[\bar{\mu}_k]], [[0]]$
  - 5:  $\mathcal{C}, \mathcal{T}$  execute Protocol 2 element-wise on inputs  $[[a_k]], [[b_k]]$ :  $\mathcal{T}$  gets  $t_k$  ▷ Secure comparison protocol
  - 6:  $\mathcal{C}, \mathcal{T}$  execute Protocol 3:  $\mathcal{C}$  obtains  $[[\mu_{k+1}]]$  ▷ Secure update protocol that ensures  $\mu_{k+1} = \max\{\bar{\mu}_k, 0\}$
- 

Line 3 ensures that the updated iterate has the required number of bits for the comparison protocol. This step is achieved by an exchange between the cloud and target node: the cloud additively blinds the iterate by a random number, sends it to the target node, which decrypts and truncates the sum and sends it back, where the cloud then subtracts the truncated random number.

The proof of security in the semi-honest model follows similar steps as in the argmax protocol in [60] and we will address it in Appendix C-A.

**Proposition 2.** *Protocol 4 is secure in the semi-honest model, according to Definition 4.*

Using the building blocks described above, we can finally assemble the protocol that privately solves the constrained quadratic optimization problem (8) with private data and sends the optimal solution to the target node. The public key  $pk_{\mathcal{T}}$  and bit-length  $l$  are known by all the parties, hence we omit them from the inputs.

---

#### PROTOCOL 5: Privacy preserving algorithm for solving strictly-convex quadratic optimization problems

**Input:**  $\mathcal{A}_{i=1,\dots,p}, b_A = \{b_j\}_{j=1,\dots,m}, c_A = \{c_j\}_{j=1,\dots,n}$ ;  $\mathcal{C}$ :  $A_C \in \mathbb{R}^{m \times n}, Q_C \in \mathbb{S}_{++}^n, K$ ;  $\mathcal{T}$ :  $sk_{\mathcal{T}}, K$

**Output:**  $\mathcal{T}$ :  $x^*$

- 1: **for**  $i=1, \dots, p$  **do**
  - 2:  $\mathcal{A}_i$ : encrypt the private information  $msg_i \leftarrow ([[b_i]], [[c_i]])$
  - 3:  $\mathcal{A}_i$ : send the encrypted messages to  $\mathcal{C}$
  - 4: **end for**
  - 5:  $\mathcal{C}$ : Construct the vectors  $[[b_A]]$  and  $[[c_A]]$  from the messages
  - 6:  $\mathcal{C}$ :  $\eta \leftarrow 1/\lambda_{\max}(A_C Q_C^{-1} A_C^T)$
  - 7:  $\mathcal{C}$ : Choose a random positive initial value  $\mu_0$  for the dual variable and encrypt it:  $[[\mu_0]]$
  - 8: **for** each  $k = 0, \dots, K - 1$  **do**
  - 9:  $\mathcal{C}, \mathcal{T}$  execute Protocol 4:  $\mathcal{C}$  gets  $[[\mu_{k+1}]]$  ▷  $\mathcal{C}, \mathcal{T}$  securely effectuate an iteration of the dual projected gradient ascent
  - 10: **end for**
  - 11:  $\mathcal{C}$ :  $[[x^*]] \leftarrow (-Q_C^{-1} A_C^T) \otimes [[\mu_K]] \oplus (-Q_C^{-1}) \otimes [[c_A]]$  and send it to  $\mathcal{T}$  ▷ Compute the primal optimum from the optimal dual solution as in (15)
  - 12:  $\mathcal{T}$ : Decrypt  $[[x^*]]$  and output  $x^*$
-

### E. Fixed-point arithmetic

The optimization problem (8) is defined on real variables, whereas the Paillier encryption scheme is defined on integers. To address this issue, we adopt a fixed-point arithmetic setting, where we allow for a number to have a fixed number of fractional bits. First, we consider numbers that have the magnitude between  $-2^{l_i-1} < x < 2^{l_i-1}$ . Second, we consider a value having  $l_i$  bits for the integer part and  $l_f$  bits for the fractional part. Therefore, by multiplying the real values by  $2^{l_f}$  and truncating the result, we obtain integers. We choose  $l = l_i + l_f$  large enough such that the loss in accuracy is negligible and assume that there is no overflow. For ease of exposition, we consider this data processing done implicitly in the protocols described.

The random numbers used for blinding the sensitive values (in Protocols 2 and 3) are sampled uniformly from the integers in  $\mathbb{Z}_N$  of  $l + \lambda$  bits, where  $\lambda$  is the statistical security parameter, as already explained in Section II, chosen such that brute-forcing the solution is intractable. In order to guarantee correctness of the comparison protocol, no overflow must take place, so we must impose  $\log_2 N > l + \lambda + 1$ .

The errors in the solution caused by the fixed-point arithmetic operations necessary for the encryption can be analyzed with the same tools as in [25], [65], [66]. The round-off errors can be regarded as states in a stable dynamical system with bounded disturbances, and hence, have a bounded norm that offers a guide on how to choose the number of fractional bits  $l_f$  for the fixed-point representation. On the other hand, the overflow and quantization errors depend on the magnitude of the dual iterates. We considered feasible and bounded problems – the dual problem (13) has a finite solution – therefore, one can select the number of integer bits  $l_i$  in the representation such that no overflow occurs.

## VI. PRIVACY OF QUADRATIC OPTIMIZATION PROTOCOL

We will now introduce the main theoretical results of the paper. In the interest of completeness of the paper, more theoretical preliminaries needed for the privacy results, such as semantic security, are given in Appendix A. We provide the statements of the results and the main ideas of the proofs here, and include the detailed arguments in Appendix C.

**Theorem 1.** *Protocol 5 achieves privacy with respect to Definition 4 for non-colluding parties.*

The intuition for the proof is as follows. Consider an iteration of the gradient ascent in Protocol 5. Firstly, in the Paillier cryptosystem, two ciphertexts are computationally indistinguishable to a party that does not have access to the decryption key. Secondly, the exchanges between the cloud and the target are additively blinded using a different random number uniformly sampled from a large enough range (at least  $\lambda$  bits more over the values that are desired to be blinded, where the size of  $\lambda$  is chosen appropriately, as discussed in Section IV-B. This means that the blinded message is statistically indistinguishable from a random number sampled from the same distribution. Thirdly, the ciphertexts are refreshed (a different encryption of the same value), after each exchange,

so a party that does not have access to the decryption key cannot infer information about the encrypted values by simply comparing the ciphertexts. Then, none of the parties can infer the magnitude or the sign of the private variables. However, we need to show that privacy is not broken by running an iteration multiple times. We prove that storing the exchanged messages does not give any new information on the private data, using similar arguments. Formally, using Definition 4, we construct a probabilistic polynomial-time simulator that randomly generates messages from the inputs and outputs such that its view and the view of the adversary, on the same inputs, are computationally indistinguishable. The correctness of Protocol 5 is immediate and follows from the correctness of the dual gradient ascent algorithm and the correctness of the comparison protocol. The detailed proof is given in Appendix C-A.

Let us now consider collusions between the parties in the setup and prove privacy of Protocol 5 under coalitions. The definition of privacy is naturally extended in this case and we can further establish the computational security in the following result – see Appendix A for more details. Definition 5 states that even under collusions, the protocol securely computes a functionality, which in this case, is solving a quadratic optimization problem. No further information is revealed than what can be inferred from the coalition’s inputs and outputs. However, if all agents and the cloud collude then they have access to all the information to solve the problem, in which case the above result is rather vacuous. Similarly, if the cloud colludes with the target node, then it can gain access to all the private data of the agents. Hence, we consider coalitions that involve either all the agents or the cloud and the target node to be prohibited (and unrealistic) and only consider coalitions between a strict subset of the agents and the cloud or between a strict subset of the agents and the target node. We give a complementary analysis on the information gained from the inputs and outputs of coalitions in Section VIII.

**Theorem 2.** *Protocol 5 achieves privacy with respect to Definition 5 against coalitions.*

The proof of Theorem 2 can be derived from the proof of Theorem 1 because the agents only contribute with their inputs to the view of the coalition, and not with new messages than. The proof is given in Appendix C-B.

## VII. ALTERNATIVE QUADRATIC OPTIMIZATION PROTOCOL WITHOUT SECURE COMPARISON

As explained in Section V-B, the major computation and communication overhead of the above protocol is the secure comparison protocol, required to project dual variables to non-negative numbers. In this section, we describe a computationally less involved alternative approach, which bypasses the need for the secure comparison protocol, at the expense of revealing more information. This approach is developed in more detail in previous work [31].

Specifically, consider a step of the dual gradient algorithm where the cloud maintains an unprojected gradient ascent step  $\bar{\mu}_k = \mu_k - \eta[A_C Q_C^{-1}(A_C^T \mu_k + c_A) + b_A]$  encrypted using the

public key of the target node. Suppose the cloud multiplies the elements in this vector with random scalar values  $r_k$  uniformly distributed over the positive integers and sends the products to the target node. The latter can decrypt the message using its private key and gain access not to the actual unprojected iterate but to the randomly scaled version of it, which reveals the sign. It can then project it to the non-negative orthant in an unencrypted fashion  $\max\{0, (r_k)_i(\bar{\mu}_k)_i\}$  for  $i = 1, \dots, m$ , and finally encrypt the result using its own public key and return it to the cloud. The cloud can divide the result with the previously selected values  $r_k$  to compute in an encrypted fashion the elements  $[[\mu_{k+1}]_i] = (1/(r_k)_i) \otimes [[\max\{0, (r_k)_i(\bar{\mu}_k)_i\}]]$  which is equivalent to the encrypted actual projected step  $[[\max\{0, (\bar{\mu}_k)_i\}]]$  for  $i = 1, \dots, m$ , because division with a positive number commutes with the max operator.

This protocol is presented next, and it can be employed as an alternative for the more demanding Protocol 4, as it bypasses the complexity of the secure comparison part. On the other hand, as can be seen from the outputs of the protocol, it reveals more information to the target node than the secure comparison approach. In particular it reveals a scaled version of the unprojected dual variable which in turn does not reveal the magnitude of this value but reveals whether this is positive, negative, or zero.

---

PROTOCOL 6: Alternative iteration of the dual projected gradient ascent method

**Input:**  $C: A_C \in \mathbb{R}^{m \times n}, Q_C \in \mathbb{S}_{++}^n, [[b_A]], [[c_A]], \eta > 0, [[\mu_k]]$ ;  $\mathcal{T}: sk_{\mathcal{T}}$

**Output:**  $C: [[\mu_{k+1}]]$ ;  $\mathcal{T}: (r_k)_i(\bar{\mu}_k)_i, i = 1, \dots, m$

- 1:  $C: [[\nabla g(\mu_k)]] \leftarrow (-A_C Q_C^{-1} A_C^T) \otimes [[\mu_k]] \oplus (-A_C Q_C^{-1}) \otimes [[c_A]] \oplus (-1) \otimes [[b_A]]$   $\triangleright$  Compute the encrypted gradient as in (14)
- 2:  $C: [[\bar{\mu}_k]] \leftarrow [[\mu_k]] \oplus \eta \otimes [[\nabla g(\mu_k)]]$   $\triangleright$  Update the value in the ascent direction
- 3:  $C$ : generate a random uniform positive scalar  $r_k$  and send  $(r_k)_i [[(\bar{\mu}_k)_i]]$  to  $\mathcal{T}$
- 4:  $\mathcal{T}$ : decrypt and truncate  $(r_k)_i \otimes [[(\bar{\mu}_k)_i]], i = 1, \dots, m$
- 5:  $\mathcal{T}$ : compute  $\max\{(r_k)_i(\bar{\mu}_k)_i, 0\}$
- 6:  $\mathcal{T}$ : encrypt  $[[\max\{(r_k)_i(\bar{\mu}_k)_i, 0\}]]$  and send to  $\mathcal{T}$
- 7:  $C$ : compute  $[[\mu_{k+1}]] = (1/(r_k)_i) \otimes [[\max\{(r_k)_i(\bar{\mu}_k)_i, 0\}]], i = 1, \dots, m$

---

Recall that we can only perform operations on integers when using the Paillier cryptosystem. To this end, we performed an implicit multiplication by  $2^{l_f}$  and truncation when encrypting numbers and a division by the corresponding factor when decrypting. For Algorithm 6 to work, we slightly modify this technique. The reasoning is the following: we want the values of  $r_k$  to be large enough to blind the magnitude of the elements of  $\bar{\mu}_k$ , so we will randomly sample  $r_k$  from the values of  $\gamma + l$  bits, where  $\gamma$  is the security size for the multiplicative blinding. Furthermore, we need to be able to represent  $1/(r_k)_i$  on the message space and to have enough precision as to obtain the correct value of  $(\mu_{k+1})_i$  in line 7. Hence, we will perform implicit multiplications and divisions by  $2^{l'_f}$ , for  $l'_f > l + \gamma$ . This change shows that, although significantly less communication is required to perform the projection operation, all the other operations are more expensive, because we work with larger numbers.

**Theorem 3.** *Protocol 5 with the iterations as in Protocol 6 achieves privacy according to Definitions 4 and 5.*

The proof follows similar steps as in the proof of the main protocol in Section VI and is sketched in Appendix C-C.

## VIII. PRIVACY DISCUSSION

In this section we discuss privacy aspects that differ from the computational notions of Definitions 4 and 5. Even if a protocol does not leak any information, the known information in a coalition (e.g., the output and some of the inputs) can be used by the coalition to infer the rest of the private inputs. More specifically, in our optimization problem, the private variables and the optimal solution are coupled via the optimality conditions (15)-(17), which are public knowledge, irrespective of the protocol, and may be used for the purpose described above.

Consider the following definition that concerns the retrieval of private data from adversarial/known data.

**Definition 6.** *(Non-unique retrieval) Let  $p$  be the private inputs of a problem and let an algorithm  $A(p)$  solve that problem. Let  $\mathcal{K}$  be the adversarial knowledge of the problem, which can contain public information, some private information (including some parts of the input) and the output of algorithm  $A$  for the adversary, denoted by  $A^{\mathcal{K}}(p)$ . We say  $p$  cannot be uniquely retrieved by the adversary if there exists a set  $\mathcal{U}$ , such that  $p \in \mathcal{U}$ ,  $|\mathcal{U}| \geq 2$  and:*

$$\forall p' \in \mathcal{U} : A^{\mathcal{K}}(p) = A^{\mathcal{K}}(p').$$

Definition 6 can be modified to be stronger by requiring the set  $\mathcal{U}$  to have an infinite number of elements.

In what follows, beyond the computational security analysis of the previous sections, we carry out an algebraic analysis on a black-box protocol that given the agent's private data  $b_A, c_A$  and the cloud's matrices  $Q_C, A_C$ , outputs the solution  $x^*$  of Problem (8) to the target node. We provide conditions such that a coalition cannot uniquely determine unknown private inputs from the output and a set of inputs, in the sense of Definition 6. In particular, this analysis applies to Protocol 5 which, assuming it runs for sufficient iterations, outputs the desired result  $x^*$  to the target node. We perform this algebraic analysis in the space of real numbers, which can be further expanded to fixed-point arithmetics for large enough precision.

Suppose without loss of generality that a coalition between  $\bar{p}$  agents ( $1 \leq \bar{p} < p$ ) has access to the elements  $b_1, \dots, b_{\bar{m}}$  with  $0 \leq \bar{m} \leq m$ , and  $c_1, \dots, c_{\bar{n}}$  with  $0 \leq \bar{n} \leq n$ . Then let us define the decomposition of the matrix  $A_C$  as:

$$A_C = \left[ \begin{array}{c|c} A_1 & \\ \hline A_{21} & A_{22} \end{array} \right] \quad (19)$$

where the matrices  $A_1 \in \mathbb{R}^{\bar{m} \times n}$ ,  $A_{21} \in \mathbb{R}^{(m-\bar{m}) \times n}$ ,  $A_{22} \in \mathbb{R}^{(m-\bar{m}) \times (n-\bar{n})}$ .

**Proposition 3.** *Consider a protocol solving Problem (8) and a coalition between the target node and agents with access to  $\bar{m}$  of the values of  $b_A$  and  $\bar{n}$  of the values of  $c_A$ . Suppose the cost and constraint matrices  $A_C, Q_C$  are public. Then:*

- (1) if  $\bar{m} < m$  and there exists a vector  $\delta \in \mathbb{R}^{m-\bar{m}}$  such that  $\delta \neq 0, \delta \geq 0$  and  $A_{21}^T \delta = 0$ , then the coalition cannot uniquely retrieve the value of  $b_A$ ;
- (2) if additionally  $\bar{n} < n$  and  $A_{22}^T \delta \neq 0$  then the coalition cannot uniquely retrieve the value of  $c_A$ .

The proof is based on the fact that the variables  $b_A, c_A, x^*, \mu^*$  satisfy the optimality conditions (15)–(17). Specifically, these are conditions that the unknown variables  $b_{\bar{m}+1}, \dots, b_m$  and  $c_{\bar{n}+1}, \dots, c_n$ , as well as the optimal dual variables  $\mu^*$ , must satisfy given all the rest of the known variables. Hence, if there are multiple such solutions to the KKT conditions the coalition cannot uniquely determine the private variables. The proof is given in Appendix C-D.

**Proposition 4.** *Consider a protocol solving Problem (8) and a coalition between the cloud and agents with access to  $\bar{m}$  of the values of  $b_A$  and  $\bar{n}$  of the values of  $c_A$ . Then, a coalition that satisfies  $\bar{m} < m$  and  $\bar{n} < n$  cannot uniquely retrieve the values of  $b_A, c_A$  and  $x^*$ .*

The proof easily follows from the fact that  $Q_C$  is a positive definite matrix, hence, the solution  $x^*$  is finite, and  $A_C$  does not have any columns or rows of zeros. A coalition between the cloud and the  $\bar{p} < p$  agents cannot solve (8) since it lacks all the data to define it ( $\bar{m} < m$  and  $\bar{n} < n$ ), so it cannot uniquely retrieve  $x^*$  and the rest of the agents' private data.

Propositions 3 and 4 give general sufficient conditions for the desired non-unique retrieval property to hold and are independent of the exact problem instance values. If the conditions stated are not satisfied, then the previous result in Theorem 2 still holds. However, the additional non-unique retrieval property may fail because the inputs and outputs of the coalition are sufficient to leak information about the rest of the private inputs. The above analysis can also be extended to the case where some of the matrices  $Q_C, A_C$  are private information.

Let us now use Definition 6 to analyze Protocol 6 and the effect the release of more information to the target node has on the privacy of the inputs of the honest agents. We perform the analysis at a setup where the protocol runs for a sufficient number of iterations and hence the dual variable  $\mu_k$  has converged to the true optimal value  $\mu^*$  and the algorithm has also converged to the true optimal primal value  $x^*$ . Suppose the target node has access to the sign of the unprojected optimal dual variable  $\mu^* + \nabla g(\mu^*)$  – note that this is the case when we employ Protocol 6. In combination with the solution  $x^*$ , this information can be further employed by the target node to infer private information of the agents.

**Proposition 5.** *Consider a protocol solving Problem (8) where the target node has access to the solution  $x^*$  and also the sign of the unprojected optimal dual variables  $t := \text{sign}(\mu^* + \nabla g(\mu^*)) \in \{+1, -1, 0\}^m$ . Suppose further the matrix  $A_C$  is publicly available. Then the private values  $b_A$  cannot be uniquely retrieved by the target node if and only if  $t_i < 0$  for some  $i \in \{1, \dots, m\}$ .*

Here  $t_i < 0$  implies that the corresponding optimal dual value is zero,  $\mu_i^* = 0$ . This means that the corresponding

constraint in problem (8) is inactive at the optimal solution  $x^*$  [61, Ch. 5] and the target node cannot uniquely determine the  $i$ th element of the corresponding bound  $b_A$ . In the opposite case, if all constraints (16) are either active or redundant at the optimal solution ( $\mu^* \geq 0$ ), this is revealed to the target node because in that case  $t_i \geq 0$ , and the private value  $b_A$  is uniquely determined by  $b_A = A_C x^*$ .

**Proposition 6.** *Consider the setup of Proposition 5 and the matrix  $Q_C$  is publicly available. The private values  $c_A$  cannot be uniquely retrieved from the outputs  $x^*$  and  $t$  of the target node if and only if  $t_i > 0$  for some  $i \in \{1, \dots, m\}$ .*

The case  $t_i > 0$  corresponds now to the case where the corresponding constraint is active at the optimal solution  $x^*$ . When this fails, all constraints are inactive at  $x^*$  and they do not play a role. Hence, we have an unconstrained quadratic problem in (8), i.e., the optimal solution satisfies the first order condition  $Q_C x^* + c_A = 0$ , which reveals the value of  $c_A$  to the target. To guarantee privacy with respect to both private values  $b_A, c_A$  we need both an inactive constraint ( $t_i < 0$ ) as well as an active constraint ( $t_j > 0$ ) for some  $i, j \in \{1, \dots, m\}$ , leaving some uncertainty in the estimations performed by the target node. Similar analysis may be performed for collusions, e.g., of the target node and some agents.

Finally, we note that here we analyzed the non-unique retrieval properties of the problem considering black-box protocols. It is of value to also consider additional steps when designing a protocol that solves Problem (8) to prevent information leakage if the conditions in the propositions are not satisfied. It may be possible to use techniques such as Differential Privacy to augment the cryptographic privacy of the protocol with output privacy, by adding correlated data to the inputs. This approach lies out of the scope of our paper, but works such as [44]–[46] may be helpful to this end.

## IX. IMPLEMENTATION

The efficiency of a secure multi-party protocol is measured in the complexity of the computations performed by each party, along with the rounds of communications between the parties. While the former is relevant from the perspective of the level of computational power required, the latter are relevant when communication is slow.

Let  $\sigma$  be the bit-length of the modulus  $N$ . Then, the size of a Paillier ciphertext is  $2\sigma$ , regardless of the size of the plaintext in  $\mathbb{Z}_N$ . A Paillier encryption takes one online exponentiation and one multiplication modulo  $N^2$ , and a decryption takes one exponentiation modulo  $N^2$  and one multiplication modulo  $N$ . The addition of two ciphertexts takes one modular multiplication modulo  $N^2$ . A multiplication of a ciphertext with a plaintext of  $l$  bits is achieved as a modular exponentiation modulo  $N^2$ . A multiplication of two elements in  $\mathbb{Z}_{N^2}^*$  can be achieved in  $O((2\sigma)^{1.7})$ . A modular exponentiation in  $N^2$  with an  $l$ -bit exponent can be computed in  $lO((2\sigma)^2)$  and can be sped up via the Chinese Remainder Theorem when the factorization of  $N$  is known. A DGK encryption takes one modular exponentiation and one modular multiplication in  $\mathbb{Z}_{N_{DGK}}^*$ , and a DGK decryption – checking if the encrypted message is 0 or not – takes an exponentiation modulo  $p_{DGK}$ .

In the setup we considered, the agents are low-power platforms and they are only required to effectuate one encryption and send one message, but the cloud and the target node are platforms with higher computational capabilities. Protocol 5 involves  $\mathcal{O}(K)$  matrix-encrypted vector multiplications and  $\mathcal{O}(Km)$  encrypted scalar additions performed at the cloud,  $\mathcal{O}(Kml)$  DGK decryptions,  $\mathcal{O}(Km)$  Paillier decryptions and  $\mathcal{O}(Km)$  encrypted scalar additions and scalar-encrypted scalar multiplications at the target node, with  $\mathcal{O}(K)$  batches of messages communicated between the two. The random numbers and their encryptions can be precomputed.

We implemented the protocol proposed in Section V-D in Python 3 and ran it on a 2.2 GHz Intel Core i7 processor. For random instances of the data in Problem (8), with  $\sigma = 1024$  and  $l = 32$  bits, and 0, 10 and 20 ms delay for communication, we obtain the average running times depicted in Figure 2 and Figure 3. The simulations are run for 30 iterations, for ease of comparison.

Figure 2 depicts the online execution time of Protocol 5 with the iterations as in Protocol 4. As expected, because we work in the dual space, the time required for running the iterations varies very little with the number of variables, but increases with the number of constraints. However, in the case of Figure 3, which depicts the online execution time of Protocol 5 with the iterations as in Protocol 6, the time varies more with the number of variables  $n$ . The reason is that we work with substantially larger numbers than in the previous case, due to the large factor with which we have to multiply in order to be able to de-mask the dual iterate, which amplifies the time difference between problems of different dimensions.

The trade-off between the privacy and communication can be seen when we artificially add a communication delay of 10 ms, respectively 20 ms between the cloud and the target node, to simulate the delays that can occur on a communication network. It can be observed in Figures 2 and 3 that the communication delay has a much smaller effect on the less private Protocol 6, which does not use the secure comparison protocol, than on the fully private Protocol 4.

#### ACKNOWLEDGMENT

This work was partially sponsored by the NSF awards 1739816, 1705135, CNS-1505799 and the Intel-NSF Partnership for Cyber-Physical Security and Privacy, ONR N00014-17-1-2012, and by TerraSwarm, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of NSF, or the U.S. Government. Andreea Alexandru would like to thank Brett Hemenway for helpful discussions.

#### REFERENCES

[1] S. Subashini and V. Kavitha, "A survey on security issues in service delivery models of cloud computing," *Journal of Network and Computer Applications*, vol. 34, no. 1, pp. 1–11, 2011.




Fig. 2. Comparison between the average running time of Protocol 5 with the iterations as in Protocol 4 for problem instances with the number of variables on the abscissa and the number of constraints in the legend. The full lines correspond to the running times of the protocol with no communication delay, the dashed lines correspond to a 10 ms communication delay and the dotted lines to a 20 ms communication delay. The simulation is run for 30 iterations, a 1024 bit key and 32 bit messages, with 16 bit precision. The statistical parameter for additive blinding is 100 bits.




Fig. 3. Comparison between the average running time of Protocol 5 with the iterations as in Protocol 6 for problem instances with the number of variables on the abscissa and the number of constraints in the legend. The full lines correspond to the running times of the protocol with no communication delay, the dashed lines correspond to a 10 ms communication delay and the dotted lines to a 20 ms communication delay. The simulation is run for 30 iterations, a 1024 bit key and 32 bit messages, with 16 bit precision. The statistical parameter for multiplicative blinding is 40 bits.

[2] D. A. Fernandes, L. F. Soares, J. V. Gomes, M. M. Freire, and P. R. Inácio, "Security issues in cloud environments: a survey," *International Journal of Information Security*, vol. 13, no. 2, pp. 113–170, 2014.

[3] "List of data breaches," [http://wikipedia.org/wiki/List\\_of\\_data\\_breaches](http://wikipedia.org/wiki/List_of_data_breaches).

[4] A. C. Yao, "Protocols for secure computations," in *Proc. 23rd Symposium on Foundations of Computer Science*. IEEE, 1982, pp. 160–164.

[5] S. Goldwasser, S. Micali, and A. Wigderson, "How to play any mental game, or a completeness theorem for protocols with an honest majority," in *Proc. 19th ACM STOC*, vol. 87, 1987, pp. 218–229.

[6] M. Bellare, V. T. Hoang, and P. Rogaway, "Foundations of garbled circuits," in *Proc. Conference on Computer and Communications Security*. ACM, 2012, pp. 784–796.

[7] R. L. Rivest, L. Adleman, and M. L. Dertouzos, "On data banks and privacy homomorphisms," *Foundations of secure computation*, vol. 4, no. 11, pp. 169–180, 1978.

[8] C. Gentry, "A fully homomorphic encryption scheme," Ph.D. dissertation, Department of Computer Science, Stanford University, 2009.

- [9] C. Gentry, A. Sahai, and B. Waters, "Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based," in *Advances in Cryptology—CRYPTO 2013*. Springer, 2013, pp. 75–92.
- [10] M. Van Dijk, C. Gentry, S. Halevi, and V. Vaikuntanathan, "Fully homomorphic encryption over the integers," in *Proc. Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2010, pp. 24–43.
- [11] Z. Brakerski and V. Vaikuntanathan, "Efficient fully homomorphic encryption from (standard) LWE," *SIAM Journal on Computing*, vol. 43, no. 2, pp. 831–871, 2014.
- [12] J. H. Cheon, A. Kim, M. Kim, and Y. Song, "Homomorphic encryption for arithmetic of approximate numbers," in *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2017, pp. 409–437.
- [13] S. Halevi, "Homomorphic encryption," in *Tutorials on the Foundations of Cryptography*. Springer, 2017, pp. 219–276.
- [14] Z. Brakerski and V. Vaikuntanathan, "Fully homomorphic encryption from ring-LWE and security for key dependent messages," in *Proc. Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2011, pp. 505–524.
- [15] T. El Gamal, "A public key cryptosystem and a signature scheme based on discrete logarithms," in *Workshop on the Theory and Application of Cryptographic Techniques*. Springer, 1984, pp. 10–18.
- [16] P. Paillier, "Public-key cryptosystems based on composite degree residuosity classes," in *Proc. Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 1999, pp. 223–238.
- [17] S. Goldwasser and S. Micali, "Probabilistic encryption & how to play mental poker keeping secret all partial information," in *Proc. 14th ACM Symposium on Theory of Computing*. ACM, 1982, pp. 365–377.
- [18] I. Damgård, M. Geisler, and M. Krøigaard, "Efficient and secure comparison for on-line auctions," in *Proc. Australasian Conference on Information Security and Privacy*. Springer, 2007, pp. 416–430.
- [19] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Proc. Theory of Cryptography Conference*. Springer, 2006, pp. 265–284.
- [20] C. Dwork, "Differential privacy: A survey of results," in *Proc. International Conference on Theory and Applications of Models of Computation*. Springer, 2008, pp. 1–19.
- [21] C. Dwork, A. Roth *et al.*, "The algorithmic foundations of differential privacy," *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.
- [22] K. Kogiso and T. Fujita, "Cyber-security enhancement of networked control systems using homomorphic encryption," in *Proc. 54th Conference on Decision and Control*. IEEE, 2015, pp. 6836–6843.
- [23] F. Farokhi, I. Shames, and N. Batterham, "Secure and private control using semi-homomorphic encryption," *Control Engineering Practice*, vol. 67, pp. 13–20, 2017.
- [24] M. S. Darup, A. Redder, I. Shames, F. Farokhi, and D. Quevedo, "Towards encrypted MPC for linear constrained systems," *IEEE Control Systems Letters*, vol. 2, no. 2, pp. 195–200, 2018.
- [25] A. B. Alexandru, M. Morari, and G. J. Pappas, "Cloud-based MPC with encrypted data," in *Proc. 57th Conference on Decision and Control*. IEEE, 2018, pp. 5014–5019.
- [26] J. Kim, C. Lee, H. Shim, J. H. Cheon, A. Kim, M. Kim, and Y. Song, "Encrypting controller using fully homomorphic encryption for security of cyber-physical systems," *IFAC-PapersOnLine*, vol. 49, no. 22, pp. 175–180, 2016.
- [27] A. Alanwar, Y. Shoukry, S. Chakraborty, P. Martin, P. Tabuada, and M. Srivastava, "PrOLoc: resilient localization with private observers using partial homomorphic encryption," in *Proc. 16th ACM/IEEE International Conference on Information Processing in Sensor Networks*. ACM, 2017, pp. 41–52.
- [28] S. Han, W. K. Ng, L. Wan, and V. C. Lee, "Privacy-preserving gradient-descent methods," *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 6, pp. 884–899, 2010.
- [29] S. Hardy, W. Henecka, H. Ivey-Law, R. Nock, G. Patrini, G. Smith, and B. Thorne, "Private federated learning on vertically partitioned data via entity resolution and additively homomorphic encryption," *arXiv preprint arXiv:1711.10677*, 2017.
- [30] C. Zhang and Y. Wang, "Privacy-preserving decentralized optimization based on ADMM," *arXiv preprint arXiv:1707.04338*, 2017.
- [31] Y. Shoukry, K. Gatsis, A. Alanwar, G. J. Pappas, S. A. Seshia, M. Srivastava, and P. Tabuada, "Privacy-aware quadratic optimization using partially homomorphic encryption," in *Proc. 55th Conference on Decision and Control*. IEEE, 2016, pp. 5053–5058.
- [32] A. B. Alexandru, K. Gatsis, and G. J. Pappas, "Privacy preserving cloud-based quadratic optimization," in *Proc. 55th Allerton Conference on Communication, Control, and Computing*. IEEE, 2017, pp. 1168–1175.
- [33] Y. Lu and M. Zhu, "Privacy preserving distributed optimization using homomorphic encryption," *arXiv preprint arXiv:1805.00572*, 2018.
- [34] W. Zheng, R. A. Popa, J. E. Gonzalez, and I. Stoica, "Helen: Maliciously secure cooperative learning for linear models," *arXiv preprint arXiv:1907.07212*, 2019.
- [35] J. Li and M. J. Atallah, "Secure and private collaborative linear programming," in *Proc. International Conference on Collaborative Computing: Networking, Applications and Worksharing*. IEEE, 2006, pp. 1–8.
- [36] W. Xie, Y. Wang, S. M. Boker, and D. E. Brown, "Privlogit: Efficient privacy-preserving logistic regression by tailoring numerical optimizers," *arXiv preprint arXiv:1611.01170*, 2016.
- [37] P. Mohassel and Y. Zhang, "SecureML: A system for scalable privacy-preserving machine learning," *Cryptology ePrint Archive*, Report 2017/396, 2017, <http://eprint.iacr.org/2017/396>.
- [38] A. Gascón, P. Schoppmann, B. Balle, M. Raykova, J. Doerner, S. Zahur, and D. Evans, "Privacy-preserving distributed linear regression on high-dimensional data," *Proceedings on Privacy Enhancing Technologies*, vol. 2017, no. 4, pp. 345–364, 2017.
- [39] M. Chase, R. Gilad-Bachrach, K. Laine, K. Lauter, and P. Rindal, "Private collaborative neural network learning," *Cryptology ePrint Archive*, Report 2017/762, 2017, <https://eprint.iacr.org/2017/762>.
- [40] J. Le Ny and G. J. Pappas, "Differentially private filtering," *IEEE Transactions on Automatic Control*, vol. 59, no. 2, pp. 341–354, 2014.
- [41] R. Dong, W. Krichene, A. M. Bayen, and S. S. Sastry, "Differential privacy of populations in routing games," in *Proc. 54th Conference on Decision and Control*. IEEE, 2015, pp. 2798–2803.
- [42] J. Cortés, G. E. Dullerud, S. Han, J. Le Ny, S. Mitra, and G. J. Pappas, "Differential privacy in control and network systems," in *Proc. 55th Conference on Decision and Control*. IEEE, 2016, pp. 4252–4272.
- [43] Y. Wang, Z. Huang, S. Mitra, and G. E. Dullerud, "Differential privacy in linear distributed control systems: Entropy minimizing mechanisms and performance tradeoffs," *IEEE Transactions on Control of Network Systems*, vol. 4, no. 1, pp. 118–130, 2017.
- [44] Z. Huang, S. Mitra, and N. Vaidya, "Differentially private distributed optimization," in *Proc. International Conference on Distributed Computing and Networking*. ACM, 2015, pp. 1–10.
- [45] 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, 2017.
- [46] 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, 2018.
- [47] J. Vaidya, "Privacy-preserving linear programming," in *Proc. ACM Symposium on Applied Computing*. ACM, 2009, pp. 2002–2007.
- [48] P. C. Weeraddana, G. Athanasiou, C. Fischione, and J. S. Baras, "Per-se privacy preserving solution methods based on optimization," in *Proc. 52nd Conference on Decision and Control*, 2013, pp. 206–211.
- [49] Y. Wang and Y. Li, "An efficient and tunable matrix-disguising method toward privacy-preserving computation," *Security and Communication Networks*, vol. 8, no. 17, pp. 3099–3110, 2015.
- [50] S. Salinas, C. Luo, W. Liao, and P. Li, "Efficient secure outsourcing of large-scale quadratic programs," in *Proc. 11th ACM on Asia Conference on Computer and Communications Security*. ACM, 2016, pp. 281–292.
- [51] S. Gade and N. H. Vaidya, "Private learning on networks," *arXiv preprint arXiv:1612.05236*, 2016.
- [52] O. Goldreich, *Foundations of Cryptography: Volume 1, Basic Tools*. Cambridge University Press, 2003.
- [53] ———, *Foundations of Cryptography: Volume 2, Basic Applications*. Cambridge University Press, 2004.
- [54] Y. Lindell, "How to simulate it – A tutorial on the simulation proof technique," in *Tutorials on the Foundations of Cryptography*. Springer International Publishing, 2017, pp. 277–346.
- [55] G. S. Vernam, "Cipher printing telegraph systems: For secret wire and radio telegraphic communications," *Journal of the AIEE*, vol. 45, no. 2, pp. 109–115, 1926.
- [56] S. M. Bellare, "Frank Miller: Inventor of the one-time pad," *Cryptologia*, vol. 35, no. 3, pp. 203–222, 2011.
- [57] Y. Dodis, "Lecture notes on exposure-resilient cryptography," 2007.
- [58] T. Veugen, "Encrypted integer division," in *Workshop on Information Forensics and Security*. IEEE, 2010, pp. 1–6.
- [59] A. Jeckmans, A. Peter, and P. Hartel, "Efficient privacy-enhanced familiarity-based recommender system," in *Proc. European Symposium on Research in Computer Security*. Springer, 2013, pp. 400–417.

- [60] R. Bost, R. A. Popa, S. Tu, and S. Goldwasser, "Machine learning classification over encrypted data," in *Proc. NDSS*, 2015.
- [61] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [62] Y. Nesterov, *Introductory lectures on convex optimization: A basic course*. Springer Science & Business Media, 2013, vol. 87.
- [63] I. Damgård, M. Geisler, and M. Krøigaard, "A correction to "Efficient and secure comparison for on-line auctions"," *International Journal of Applied Cryptography*, vol. 1, no. 4, pp. 323–324, 2009.
- [64] T. Veugen, "Improving the DGK comparison protocol," in *Workshop on Information Forensics and Security*. IEEE, 2012, pp. 49–54.
- [65] J. L. Jerez, P. J. Goulart, S. Richter, G. A. Constantinides, E. C. Kerrigan, and M. Morari, "Embedded online optimization for model predictive control at megahertz rates," *IEEE Transactions on Automatic Control*, vol. 59, no. 12, pp. 3238–3251, 2014.
- [66] M. Rubagotti, P. Patrinos, A. Guiggiani, and A. Bemporad, "Real-time model predictive control based on dual gradient projection: Theory and fixed-point fpga implementation," *International Journal of Robust and Nonlinear Control*, vol. 26, no. 15, pp. 3292–3310, 2016.
- [67] J. Garay, B. Schoenmakers, and J. Villegas, "Practical and secure solutions for integer comparison," in *International Workshop on Public Key Cryptography*. Springer, 2007, pp. 330–342.
- [68] V. Kolesnikov, A.-R. Sadeghi, and T. Schneider, "Improved garbled circuit building blocks and applications to auctions and computing minima," in *Proc. International Conference on Cryptology and Network Security*. Springer, 2009, pp. 1–20.
- [69] H. Lipmaa and T. Toft, "Secure equality and greater-than tests with sublinear online complexity," in *Proc. International Colloquium on Automata, Languages, and Programming*. Springer, 2013, pp. 645–656.
- [70] G. Couteau, "Efficient secure comparison protocols," *Cryptology ePrint Archive*, Report 2016/544, 2016, <http://eprint.iacr.org/2016/544>.

## APPENDIX A CRYPTOGRAPHIC TOOLS

Proving a protocol is secure in the semi-honest model involves a heavy use of tools from cryptography, which we present here for the interested reader. In this Appendix, we aim to give closure to the security notions introduced in Section III, as well as provide the details of the Paillier encryptions scheme, used throughout the paper and generically introduced in Section IV. Furthermore, we detail the tools used in the proof of Theorem 1, such as semantic security.

### A. Security of multi-party protocols in the semi-honest model

Definition 4 and Definition 5 state that the view of any of the parties participating in the protocol, on each possible set of inputs, can be simulated based only on its own input and output. The intuition behind these definitions can also be viewed as follows: consider a black-box protocol that receives the inputs of Problem (8) and solves it, not exchanging any messages in order to achieve the solution. Such a black-box protocol satisfies Definitions 4 and 5. In other words, meeting the requirements of the above definitions is equivalent to not revealing any other information than what is already known to any of the parties, i.e. inputs, prescribed outputs, if any, and previously known side information, meaning details about the optimization algorithm.

**Remark A.1.** ([53, Ch. 7], [52, Ch. 4]) *Auxiliary inputs*, which are inputs that capture additional information available to each of the parties (e.g. local configurations, side-information), are implicit in Definition 4 and Definition 5.

**Remark A.2.** *Satisfying Definitions 4 and 5 is a stronger requirement than guaranteeing that an adversary cannot*

*uniquely retrieve the data of the honest parties, i.e. Definition 6.*

Revealing sensitive information does not always lead to a unique retrieval of the private data. Nevertheless, any piece of information revealed by the execution of the protocol, that cannot be obtained only from its inputs and outputs, leads to the violation of Definitions 4, 5, even if the private data cannot be singled out with this information.

### B. Paillier cryptosystem

In this part, we will provide the rest of the details regarding the implementation of the Paillier cryptosystem, introduced in Section IV. The pair of keys corresponding to this cryptosystem is  $(pk, sk)$ , where the public key is  $pk = (N, g)$  and the secret key is  $sk = (\gamma, \delta)$ .  $N$  is called the modulus and is the product of two large prime numbers  $p, q$ , and  $g$  is an element of order  $N$  in the multiplicative group  $\mathbb{Z}_N^*$ . Furthermore,

$$\gamma = \text{lcm}(p-1, q-1), \quad \delta = ((g^\gamma \bmod N^2 - 1)/N)^{-1} \bmod N.$$

A simpler variant to generate the pair of keys when  $p$  and  $q$  have the same number of bits is to choose  $g = N + 1$ ,  $\gamma = \phi(N) = (p-1)(q-1)$ , where  $\phi$  is Euler's totient function and  $\delta = \phi(N)^{-1} \bmod N$ .

For a plaintext  $a \in \mathbb{Z}_N$ , the Paillier encryption is:

$$[[a]] = g^a r^N \bmod N^2, \quad (20)$$

where  $r$  is a random non-zero element in  $\mathbb{Z}_N$ , which makes Paillier a probabilistic encryption scheme. Therefore, the ciphertexts do not preserve the order relation between the plaintexts. The decryption primitive is the following:

$$a = ((([a]]^\gamma \bmod N^2 - 1)/N)\delta \bmod N, \quad (21)$$

which uses the fact that  $(1 + N)^a = 1 + Na \bmod N^2$ .

In the Paillier scheme, in order to obtain addition between plaintexts, the operation between ciphertexts is modular multiplication, which was denoted by  $\oplus$  in the text:

$$\begin{aligned} [[a]] \cdot [[b]] &= g^a r^N \bmod N^2 \cdot g^b r'^N \bmod N^2 \\ &= g^{a+b} (rr')^N \bmod N^2 \\ &= [[a + b]] \bmod N^2. \end{aligned} \quad (22)$$

Negation is achieved by modular inverse:

$$\begin{aligned} [[a]]^{-1} &= g^{-a} (r^{-1})^N \bmod N^2 \\ &= [[-a]] \bmod N^2. \end{aligned} \quad (23)$$

The multiplication between a plaintext value  $c$  and an encrypted value  $[[a]]$ , which was denoted by  $c \otimes [[a]]$  in the text, is obtained in the following way:

$$[[a]]^c = g^{ac} (r^c)^N \bmod N^2 = [[ca]]. \quad (24)$$

**Remark A.3.** *The Paillier cryptosystem is malleable, which means that a party that does not have the private key can alter a ciphertext such that another valid ciphertext is obtained. Malleability is a desirable property in order to achieve cloud-outsourced computation on encrypted data, but allows ciphertext attacks. In this work, we assume that the parties have access to authenticated channels, therefore an adversary cannot alter the messages sent by the honest parties.*

### C. Semantic security of a cryptosystem

Proving security in the semi-honest model involves the concepts of semantic security of an encryption scheme. Under the assumptions of decisional composite residuosity [16] and hardness of factoring, the partially homomorphic encryption schemes Paillier and DGK are semantically secure and have indistinguishable encryptions, which, in essence, means that an adversary that has the plaintext messages  $a$  and  $b$  cannot distinguish between the encryption  $[[a]]$  and encryption  $[[b]]$ .

**Definition A.1.** (Semantic Security [53, Ch. 5]) An encryption scheme is **semantically secure** if for every probabilistic polynomial-time algorithm,  $A$ , there exists a probabilistic polynomial-time algorithm  $A'$  such that for every two polynomially bounded functions  $f, h : \{0, 1\}^* \rightarrow \{0, 1\}^*$  and for any probability ensemble  $\{X_n\}_{n \in \mathbb{N}}$ ,  $|X_n| \leq \text{poly}(n)$ , for any positive polynomial  $p$  and sufficiently large  $n$ :

$$\begin{aligned} \Pr[A(E(X_n), h(X_n), 1^n) = f(X_n)] &< \\ \Pr[A'(h(X_n), 1^n) = f(X_n)] + 1/p(n), \end{aligned}$$

where  $E(\cdot)$  is the encryption primitive.

The security parameter  $n$  is given as input in the form  $1^n$  to  $A$  and  $A'$  to ensure they can run in  $\text{poly}(n)$ -time.

## APPENDIX B COMPARISON PROTOCOL

Consider a two-party computation problem under an encryption scheme that does not support comparison between encrypted data. A number of secure comparison protocols on private inputs owned by two parties have been proposed in the literature [18], [67]–[69] etc., with a survey of the state of the art given in [70]. Most of the comparison protocols have linear complexity in the size of the inputs, since the comparison is done bitwise. Out of these, the DGK protocol proposed in [18] with the correction in [63] remains one of the most computationally efficient protocols, and [68] is the most efficient for large-scale protocols. [69] proposes a comparison that is sublinear in the number of invocation of the cryptographic primitive, but has greater communication complexity and is only competitive for large inputs, due to the constants involved. Depending on the specific problem, some variants might be better than others. Since this comparison protocol is used as a block, it can be easily replaced in Protocol 2 once more efficient protocol are proposed.

In [18], [63], Damgård, Geisler and Krøigaard describe a protocol for secure comparison protocol and towards that functionality, they propose the DGK additively homomorphic encryption scheme with the property that it is efficient to determine if the value zero is encrypted, which is useful for comparisons and when working with bits. The authors also prove the semantic security of the DGK cryptosystem under the hardness of factoring assumption.

The plaintext space for DGK is  $\mathbb{Z}_u$ , where  $u$  is a small prime divisor of  $p - 1$  and  $q - 1$ , where  $p, q$  are large, same size prime numbers, and  $N = pq$ . Parameters  $v_p$  and  $v_q$  are  $t$ -bit prime divisors of  $p - 1$  and  $q - 1$ . The numbers  $g$  and  $h$  are elements of  $\mathbb{Z}_N^*$  of order  $uv_p v_q$  and  $v_p v_q$ . The DGK

encryption scheme has public key  $pk_{DGK} = (N, g, h, u)_{DGK}$  and  $sk_{DGK} = (p, q, v_p, v_q)_{DGK}$ . For a plaintext  $x \in \mathbb{Z}_u$ , a DGK encryption is:

$$[a] = g^a h^r \bmod N,$$

with  $r$  a random  $2t$ -bits integer, which makes DGK a probabilistic encryption scheme. We use this encryption scheme with the purpose of encrypting and decrypting bits that represent comparison results, therefore, for decryption, we only need to check if the encrypted value is 0. It is enough to verify:

$$[a]^{v_p v_q} \bmod p = 1$$

to see if  $a = 0$ . Since DGK is an additively homomorphic scheme, the following holds, where  $D(\cdot)$  is the DGK decryption primitive:

$$a + b = D([a] \cdot [b] \bmod N), \quad -a = D([a]^{-1} \bmod N). \quad (25)$$

We will use again  $\oplus$  and  $\otimes$  to abstract the addition between and encrypted values, respectively, the multiplication between a plaintext and an encrypted values.

Protocol 2 calls the DGK protocol, presented in [64, Protocol 3]. The two parties have each their own private input  $\alpha$ , respectively  $\beta$ . Using the binary representations of  $\alpha$  and  $\beta$ , the two parties exchange  $l$  blinded values such that each of the parties will obtain a bit  $\delta_C \in \{0, 1\}$  and  $\delta_T \in \{0, 1\}$  that satisfy the following relation:  $\delta_C \vee \delta_T = (\alpha \leq \beta)$  after executing Protocol 7, where  $\vee$  denotes the exclusive or operation. Then, in Protocol 2 line 8,  $\mathcal{C}$  obtains the encrypted bit  $[[t']]$  that satisfies  $(t' = 1) \Leftrightarrow (\beta < \alpha)$ .

To guarantee the security of the DGK scheme, we choose a key-size that makes factoring hard and set the randomization in the encryption primitive to be of length greater than  $2t$  bits, for  $t = 160$ . See more details in [18], [63].

---

PROTOCOL 7: Protocol for secure two-party comparison with plaintext inputs using DGK [18], [64]

**Input:**  $\mathcal{C}$ :  $\alpha$ ;  $\mathcal{T}$ :  $\beta, sk_{\mathcal{T}}, sk_{DGK}$

**Output:**  $\mathcal{C}$ :  $\delta_C$ ;  $\mathcal{T}$ :  $\delta_T$  such that  $\delta_C \oplus \delta_T = (\alpha \leq \beta)$

- 1:  $\mathcal{T}$ : send the encrypted bits  $[\beta_i]$ ,  $0 \leq i < l$  to  $\mathcal{C}$
  - 2: **for** each  $0 \leq i < l$  **do**
  - 3:  $\mathcal{C}$ :  $[\alpha_i \vee \beta_i] \leftarrow \begin{cases} [\beta_i] & \text{if } \alpha_i = 0 \\ [1] \oplus (-1) \otimes [\beta_i] & \text{otherwise} \end{cases}$
  - 4: **end for**
  - 5:  $\mathcal{C}$ : Choose uniformly random bit  $\delta_C \in \{0, 1\}$
  - 6:  $\mathcal{C}$ : Compute set  $\mathcal{L} = \{i | 0 \leq i < l \text{ and } \alpha_i = \delta_C\}$
  - 7: **for** each  $i \in \mathcal{L}$  **do**
  - 8:  $\mathcal{C}$ : compute  $[c_i] \leftarrow \sum_{j=i+1}^l [\alpha_j \vee \beta_j]$  ▷ encrypted sum
  - 9:  $\mathcal{C}$ :  $[c_i] \leftarrow \begin{cases} [1] \oplus [c_i] \oplus (-1) \otimes [\beta_i] & \text{if } \delta_C = 0 \\ [1] \oplus [c_i] & \text{otherwise} \end{cases}$
  - 10: **end for**
  - 11:  $\mathcal{C}$ : generate uniformly random non-zero values  $r_i$  of  $2t$  bits,  $0 \leq i < l$
  - 12: **for** each  $0 \leq i < l$  **do**
  - 13:  $\mathcal{C}$ :  $[c_i] \leftarrow \begin{cases} r_i \otimes [c_i] & \text{if } i \in \mathcal{L} \\ [r_i] & \text{otherwise} \end{cases}$
  - 14: **end for**
  - 15:  $\mathcal{C}$ : send the values  $[c_i]$  in random order to  $\mathcal{T}$
  - 16:  $\mathcal{T}$ : if at least one of the values  $c_i$  is decrypted to zero, set  $\delta_T \leftarrow 1$ , otherwise set  $\delta_T \leftarrow 0$
-



APPENDIX C  
PROOF OF MAIN RESULTS

A. Proof of Theorem 1

For simplicity of the exposition, we avoid mentioning the public data (public keys  $pk_{\mathcal{T}}, pk_{DGGK}$ , number of iterations  $K$  and number of bits  $l$ ) in the views, since they are public. In what follows, we will successively discuss the views of each type of party participating in the protocol: agents, cloud and target node. As mentioned in Definitions 4 and 5, the views of the parties during the execution of Protocol 5 are constructed on all the inputs of the parties involved and symbolize the real values the parties get in the execution of the protocol. We will denote by  $\mathcal{I}$  the inputs of all the parties:

$$\mathcal{I} = \{b_{\mathcal{A}}, c_{\mathcal{A}}, A_{\mathcal{C}}, Q_{\mathcal{C}}, sk_{\mathcal{T}}, sk_{DGGK}\}.$$

Furthermore, in order to construct a simulator that simulates the actions of a party, we need to feed into it the inputs and outputs of that respective party.

1) *Simulator for one agent  $\mathcal{A}_i$* : Agent  $\mathcal{A}_i$ , for every  $i = 1, \dots, p$ , has the following inputs  $I_{\mathcal{A}_i} = (\{b_j\}_{j=1, \dots, m_i}, \{c_j\}_{j=1, \dots, n_i})$ . To avoid cluttering, we will drop the subscripts  $j = 1, \dots, m_i$  and  $j = 1, \dots, n_i$ . Then agent  $\mathcal{A}_i$  has the following view:

$$V_{\mathcal{A}_i}(\mathcal{I}) := (b_j, c_j, [[b_j]], [[c_j]], \text{coins}),$$

where coins represent the random values used for encryption.

The agents are only online to send their encrypted data to the cloud, and they do not have any role and output afterwards. Hence, a simulator  $S_{\mathcal{A}_i}$  would simply generate the random values necessary to encrypt its inputs as specified by the protocol and output the view obtained such:

$$S_{\mathcal{A}_i} := (b_j, c_j, [\widetilde{b_j}], [\widetilde{c_j}], \widetilde{\text{coins}}),$$

where by  $\widetilde{(\cdot)}$  we denote the quantities obtained by the simulator, which are different than the quantities of the agents, but follow the same distribution. Hence, it is trivial to see that the protocol is secure in the semi-honest model from the point of view of the interaction with the agents.

Next, we want to prove the privacy of the protocol from the point of view of the interactions with the cloud and the target node. We will construct a sequence of algorithms such that we obtain that the view of the real parties after the execution of  $K$  iterations is the same as the view of simulators that simply execute  $K$  iterations with random exchanged messages.

2) *Simulator for the cloud  $\mathcal{C}$* : For the simplicity of the exposition, we will treat  $\mu$ ,  $b_{\mathcal{A}}$  and  $c_{\mathcal{A}}$  as scalars. The same steps are repeated for every element in the vectors. For clarity, we will form the view of the cloud in steps, using pointers to the lines in Protocol 5. The view of the cloud during the execution of lines 5-6 is:

$$V_{\mathcal{C}}^{-1}(\mathcal{I}) = (A_{\mathcal{C}}, Q_{\mathcal{C}}, \eta, [[b_{\mathcal{A}}]], [[c_{\mathcal{A}}]], [[\mu_0]], \text{coins}) =: I_{\mathcal{C}}^{-1}, \quad (26)$$

where coins represent the random values generated for the Paillier encryption. Furthermore, we construct the view of the cloud at iteration  $k = 0, \dots, K - 1$  during the execution of an instance of Protocol 4, which will be constructed on the

inputs of all parties: the inputs  $\mathcal{I}$  and the data the parties had at iteration  $k - 1$ . We denote the view of the cloud at iteration  $k - 1$  by  $I_{\mathcal{C}}^{k-1}$  which, along with  $\mathcal{I}$  and the view of the target node at iteration  $k - 1$ , denoted by  $I_{\mathcal{T}}^{k-1}$  (30), will be what the view is constructed on at iteration  $k$ .

$$\bar{\mathcal{I}}^{k-1} := \mathcal{I} \cup I_{\mathcal{C}}^{k-1} \cup I_{\mathcal{T}}^{k-1}, \quad (27)$$

$$I_{\mathcal{C}}^k := V_{\mathcal{C}}^k(\bar{\mathcal{I}}^{k-1}) = (I_{\mathcal{C}}^{k-1}, [[\mu_k]], [[\bar{\mu}_k]], \underbrace{\pi_k, \text{coins}_{1k}}_{\text{Protocol 1}}, \underbrace{\rho_k, [[t_k]], m^{\text{comp}_k}, \text{coins}_{2k}}_{\text{Protocol 2}}, \underbrace{r_k, s_k, [[v_k]], [[\mu_{k+1}]], \text{coins}_{3k}}_{\text{Protocol 3}}), \quad (28)$$

where  $m^{\text{comp}_k}$  are messages exchanged in the comparison protocol, and  $\text{coins}_{jk}$ , for  $j = 1, 2, 3$  are the random numbers generated in Protocol  $j$ . Finally, the view of the cloud after the execution of line 11 in Protocol 5 is:

$$V_{\mathcal{C}}^K(\bar{\mathcal{I}}^{K-1}) = (I_{\mathcal{C}}^{K-1}, [[x^*]]). \quad (29)$$

Therefore, the view of the cloud during the whole execution of Protocol 5 is:

$$V_{\mathcal{C}}(\mathcal{I}) := V_{\mathcal{C}}^K(\bar{\mathcal{I}}^{K-1}).$$

We first construct a simulator on the inputs  $I_{\mathcal{C}} = \{A_{\mathcal{C}}, Q_{\mathcal{C}}\}$  that mimics  $V_{\mathcal{C}}^{-1}(\mathcal{I})$  in (26):

- Generate  $n + m$  random numbers of  $l$  bits  $\widetilde{b_{\mathcal{A}}}, \widetilde{c_{\mathcal{A}}}$ ;
- Generate a random positive initial value  $\widetilde{\mu_0}$ ;
- Generate  $n + m + 1$  uniformly random numbers for the Paillier encryption and denote them  $\widetilde{\text{coins}}$ ;
- Compute  $[[\widetilde{b_{\mathcal{A}}}], [[\widetilde{c_{\mathcal{A}}}], [[\widetilde{\mu_0}]]$ ;
- Compute  $\eta$  following line 6;
- Output  $\widetilde{I_{\mathcal{C}}^{-1}} := S_{\mathcal{C}}^{-1}(I_{\mathcal{C}}) = (A_{\mathcal{C}}, Q_{\mathcal{C}}, [[\widetilde{b_{\mathcal{A}}}], [[\widetilde{c_{\mathcal{A}}}], \eta, [[\widetilde{\mu_0}]], \widetilde{\text{coins}})$ .

Since Protocol 4 is secure in the semi-honest model (Proposition 2), we know that there exists a probabilistic polynomial-time (ppt) simulator for the functionality of Protocol 4 on inputs  $(A_{\mathcal{C}}, Q_{\mathcal{C}}, [[b_{\mathcal{A}}]], [[c_{\mathcal{A}}]], \eta, [[\mu_k]])$  and output  $[[\mu_{k+1}]]$ . However, we need to show that we can simulate the functionality of consecutive calls of Protocol 4, or, equivalently, on one call of Protocol 4 but on the augmented input that contains the data of the cloud in the previous iterations. Call such a simulator  $S_{\mathcal{C}}^k$ , that on the input  $I_{\mathcal{C}}^{k-1}$  mimics  $V_{\mathcal{C}}^k(\bar{\mathcal{I}}^{k-1})$  in (28), for  $k = 0, \dots, K - 1$ :

- Compute  $[[\nabla g(\mu_k)]]$  and  $[[\bar{\mu}_k]]$  as in lines 1-2 of Protocol 4 from  $[[\mu_k]], [[b_{\mathcal{A}}]], [[c_{\mathcal{A}}]]$  which are included in  $I_{\mathcal{C}}^{k-1}$ ;
- Generate a random permutation  $\widetilde{\pi}_k$  and apply it on  $([[0]], [[\bar{\mu}_k]])$  as in Protocol 1;
- Follow Protocol 2 and replace the incoming messages by DGGK encryptions of appropriately sized random values to obtain  $\widetilde{\rho}_k, m^{\text{comp}_k}$ ;
- Generate random bit  $\widetilde{t}_k$  and its encryption  $[[\widetilde{t}_k]]$ ;
- Generate random values  $\widetilde{r}_k$  and  $\widetilde{s}_k$  as in line 1 in Protocol 3 and their encryptions  $[[\widetilde{r}_k]], [[\widetilde{s}_k]]$ ;
- Obtain  $[[\widetilde{v}_k]]$  by choosing between the elements of  $\widetilde{\pi}_k([0], [[\bar{\mu}_k]]) + (\widetilde{r}_k, \widetilde{s}_k)$  according to the generated  $\widetilde{t}_k$ ;
- Compute  $[[\widetilde{\mu}_{k+1}]]$  as in line 8 of Protocol 3;
- Denote the rest of the random values used for encryption and blinding by  $\widetilde{\text{coins}}_k$ ;

• Output  $\tilde{I}_C^k := S_C^k(I_C^{k-1}) = (I_C^{k-1}, [[\tilde{\mu}_k]], [[\tilde{\pi}_k]], [[\tilde{z}_k]], \underbrace{m^{\text{comp}_k}, [[\tilde{t}_k]], [[\tilde{r}_k]], [[\tilde{s}_k]], [[\tilde{v}_k]], [[\tilde{\mu}_{k+1}}], \text{coins}_k}_{\text{Protocol 2}})$ .

Finally, a trivial simulator  $\tilde{I}_C^K := S_C^K(I_C^{K-1})$  for  $V_C^K(\tilde{\mathcal{I}}^{K-1})$  is obtained by simply performing line 11 on the inputs.

**Proposition C.1.**  $S_C^k(I_C^{k-1}) \stackrel{c}{\equiv} V_C^k(\tilde{\mathcal{I}}^{k-1})$ , for  $k = -1, \dots, K$ , where  $I_C^{k-2} := I_C$ .

*Proof.* For  $k = -1$ , due to the fact that the coins and the initial iterate are generated by the simulator via the same distributions as specified in the protocol, and due to the semantic security of the Paillier encryption, which guarantees indistinguishability of encryptions, we obtain that  $V_C^{-1}(\mathcal{I}) \stackrel{c}{\equiv} S_C^{-1}(I_C)$ .

For  $0 \leq k \leq K-1$ ,  $(\text{coins}_{1k}, \text{coins}_{2k}, \text{coins}_{3k}) \stackrel{s}{\equiv} \underbrace{\text{coins}_k}_{\text{Protocol 2}}$  because they are generated from the same distributions. Similarly, the distributions of  $(\tilde{\pi}_k, \tilde{\pi}_k([[0]], [[\tilde{\mu}_k]]))$  and  $(\pi_k, \pi_k([[0]], [[\tilde{\mu}_k]]))$  are the same and  $\tilde{\pi}_k, \pi_k$  are independent of the other parameters. The quantities from the DGK protocol  $\tilde{\rho}_k \stackrel{s}{\equiv} \rho_k$  and  $\underbrace{m^{\text{comp}_k}}_{\text{Protocol 2}} \stackrel{c}{\equiv} m^{\text{comp}_k}$  either due to the semantic security of the DGK protocol, e.g.,  $[\tilde{t}'], [\tilde{\delta}_{\mathcal{T}}]$  or due to having the same distributions, e.g.,  $\tilde{\rho}_k, \tilde{\delta}_{\mathcal{C}}, \tilde{\alpha}$ . The values in the update step  $\tilde{r}_k, \tilde{s}_k$  are sampled from the same distribution as  $r_k, s_k$ , and, finally,  $[[\tilde{t}_k]], [[\tilde{v}_k]], [[\tilde{\mu}_{k+1}}] \stackrel{c}{\equiv} [[t_k]], [[v_k]], [[\mu_{k+1}]]$  due to the semantic security of the Paillier encryption. Thus,  $S_C^k(I_C^{k-1}) \stackrel{c}{\equiv} V_C^k(\tilde{\mathcal{I}}^{k-1})$ .

For  $k = K$ ,  $S_C^K(I_C^{K-1}) \stackrel{c}{\equiv} V_C^K(\tilde{\mathcal{I}}^{K-1})$  follows from the fact that the simulator simply executes the last part of the protocol and from the Paillier scheme's semantic security.  $\square$

Hence, we obtained that  $I_C^k \stackrel{c}{\equiv} \tilde{I}_C^k$ , for  $k = -1, \dots, K$ . The essence of the proof of Proposition C.1 is that all the messages the cloud receives are encrypted. Then, thanks to the semantic security of the Paillier and DGK schemes, the extra information included in  $\tilde{\mathcal{I}}^{k-1}$  from the previous iterations cannot be used to extract other information about the values at iteration  $k$ . From this, we also have the next corollary:

**Corollary C.1.**  $S_C^k(\tilde{I}_C^{k-1}) \stackrel{c}{\equiv} S_C^k(I_C^{k-1})$ , equivalent to  $S_C^{k+1}(S_C^k(I_C^{k-1})) \stackrel{c}{\equiv} S_C^{k+1}(V_C^k(\tilde{\mathcal{I}}^{k-1}))$ , for any  $k = 0, \dots, K-1$ .

Finally, we construct a simulator  $S_C(I_C)$  for the execution of Protocol 5 and we will show that its view will be computationally indistinguishable from  $V_C(\mathcal{I})$ . To this end, we define the following sequence of views – obtained as hybrids between the real views and the views of the simulators:

$$\begin{aligned} V_C(\mathcal{I}) &= H_{-1}(\mathcal{I}) = V_C^K(\tilde{\mathcal{I}}^{K-1}) \\ H_0(\mathcal{I}) &= S_C^K \circ V_C^{K-1}(\tilde{\mathcal{I}}^{K-2}) \\ H_1(\mathcal{I}) &= S_C^K \circ S_C^{K-1} \circ V_C^{K-2}(\tilde{\mathcal{I}}^{K-3}) \\ &\vdots \\ H_K(\mathcal{I}) &= S_C^K \circ S_C^{K-1} \circ \dots \circ S_C^0 \circ V_C^{-1}(\mathcal{I}) \\ S_C(I_C) &= H_{K+1}(\mathcal{I}) = S_C^K \circ S_C^{K-1} \circ \dots \circ S_C^0 \circ S_C^{-1}(I_C). \end{aligned}$$

By transitivity,  $H_{-1}$  and  $H_{K+1}$  are computationally indistinguishable if:

$$H_{-1} \stackrel{c}{\equiv} H_0 \stackrel{c}{\equiv} \dots \stackrel{c}{\equiv} H_k \stackrel{c}{\equiv} H_{k+1} \stackrel{c}{\equiv} H_{k+2} \stackrel{c}{\equiv} \dots \stackrel{c}{\equiv} H_{K+1}.$$

This result follows from induction on Corollary C.1. In conclusion, we obtain that  $S_C(I_C) \stackrel{c}{\equiv} V_C(\mathcal{I})$ , which verifies that Protocol 5 achieves privacy with respect to the cloud.

3) *Simulator for the target node  $\mathcal{T}$* : We proceed with the same steps in order to show that the consecutive  $K$  iterations form a protocol that is secure in the semi-honest model from the point of view of the target node. We will use the secrecy of the one-time pad variant used for blinding. The symmetric encryption used in this paper, as discussed in Section IV-B, to which we refer as blinding, guarantees that a value of  $l$  bits additively blinded by a random value of  $l + \lambda$  bits is statistically indistinguishable from a random value of  $l + \lambda + 1$  bits.

The inputs and output of the target node in Protocol 5 are  $I_{\mathcal{T}} = (sk_{\mathcal{T}}, sk_{DGK}, x^*)$ . As in (27),  $\tilde{\mathcal{I}}^{k-1}$  represents the inputs of all the parties at iteration  $k$ , with  $\tilde{\mathcal{I}}^{-1} = \mathcal{I}$ . Then, the view of the target node during iteration  $k = 0, \dots, K-1$  is:

$$I_{\mathcal{T}}^k := V_{\mathcal{T}}^k(\tilde{\mathcal{I}}^{k-1}) = (sk_{\mathcal{T}}, sk_{DGK}, \underbrace{z_k, t_k, m^{\text{comp}_k}, \text{coins}_{2k}}_{\text{Protocol 2}}, \underbrace{\bar{a}_k, \bar{b}_k, v_k, \text{coins}_{3k}}_{\text{Protocol 3}}). \quad (30)$$

The view of the target node during the last step of the protocol is:

$$V_{\mathcal{T}}^K(\tilde{\mathcal{I}}^{K-1}) := (I_{\mathcal{T}}^{K-1}, [[x^*]]). \quad (31)$$

As before, the view of the target during the execution of Protocol 5 is:

$$V_{\mathcal{T}}(\mathcal{I}) := V_{\mathcal{T}}^K(\tilde{\mathcal{I}}^{K-1}).$$

In order to be able to construct a simulator with indistinguishable view from the view of the target node, we need to show that the target node is not capable of inferring new relevant information about the private data (other than what can be inferred solely from its inputs and outputs) although it knows the optimal solution  $x^*$  and has access to the messages from multiple iterations.

Apart from the last message, which is the encryption of the optimization solution  $[[x^*]]$ , and the comparison results  $t_k$ , all the values the target node receives are blinded, with different sufficiently large values at each iteration. The target node knows that  $Q_C x^* = A_C^T \mu_K - c_A$  and that  $\mu_K = v_K - \bar{r}_K$ , for some random value  $\bar{r}_K$ . However, from these two equations, although it has access to  $x^*$  and  $v_K$ , the target node cannot infer any new information about the private data of the agents or about  $\mu_K$ , even when  $Q_C, A_C$  are public, thanks to the large enough randomization.

Let  $a_k, b_k = \pi_k(0, \bar{\mu}_k)$ , where  $\pi_k$  is a random permutation. From Protocol 2, the target receives  $z_k$  which is the additively blinded value of  $b_k - a_k + 2^l$ , and other blinded values, denoted in (30) as  $m^{\text{comp}_k}$ . Hence, provided the blinding noises are refreshed at every iteration, the target node cannot infer any information, as follows from Section IV-B. At the end of the protocol, it receives the bit  $t_k$  which is 1 if  $a_k \leq b_k$  and 0 otherwise. However, due to the uniform randomization between the order of  $\bar{\mu}_k$  and 0 at every iteration for assigning the values of  $a_k$  and  $b_k$ , described in Protocol 1, the target node cannot identify the sign of  $\bar{\mu}_k$ , and by induction, the sign of  $\bar{\mu}_{K-1}$  and magnitude of  $\mu_K$ . This means that having

access to  $x^* = x_K$  does not bring more information about the blinded values from the intermediary steps.

We now investigate the relation between the messages from consecutive iterations. From the update protocol 3, we know:

$$\begin{aligned} v_k &= (a_k + r_k)(1 - t_k) + (b_k + s_k)t_k \\ \mu_{k+1} &= v_k - r_k(1 - t_k) - s_k t_k = a_k(1 - t_k) + b_k t_k, \end{aligned}$$

where the variables are vectors. The target node knows the values of  $v_k$  and  $t_k$ , but  $\pi_k, \bar{\mu}_k, r_k, s_k$  are unknown to it. Furthermore, let  $\pi_{k+1}$  be the permutation applied by the cloud at step  $k + 1$  in Protocol 1, unknown to the target node. Take for example the case when  $t_k = 0$  and  $t_{k+1} = 0$ . Let  $\tilde{Q} = \mathbf{I} - \eta A_C Q_C^{-1} A_C^T$  and  $\mathbf{E} = [\mathbf{I} \ \mathbf{0}]$ . Then:

$$\begin{aligned} v_{k+1} &= a_{k+1} + r_{k+1} = \mathbf{E} \pi_{k+1}(0, \tilde{Q}(v_k - r_k) - \eta A_C Q_C^{-1} c_A - \\ &\quad - \eta b_A) + r_{k+1}. \end{aligned}$$

The above equation shows the target node cannot construct  $v_{k+1}$  from  $v_k$ . Similar equations arise when considering the other values for  $t_k$  and  $t_{k+1}$ . This guarantees that an integer  $\tilde{v}_k^i$  obtained by selecting uniformly at random from  $(2^{l+\lambda}, 2^{l+\lambda+1}) \cap \mathbb{Z}_N$  will have the distribution statistically indistinguishable from  $v_k^i$ . Moreover:

$$\begin{aligned} \mu_{k+2} &= \max\{0, \tilde{Q}\mu_{k+1} - \eta A_C Q_C^{-1} c_A - \eta b_A\} \\ &= \max\{0, \tilde{Q}(v_k - r_k(1 - t_k) - s_k t_k) - \eta A_C Q_C^{-1} c_A - \eta b_A\} \\ &= v_{k+1} - r_{k+1}(1 - t_{k+1}) - s_{k+1} t_{k+1}. \end{aligned}$$

Since the blinding noise is different at each iteration and uniformly sampled, the target node cannot retrieve  $\mu_{k+2}$  from  $v_k$  and  $v_{k+1}$ . In short, if the target node receives some random values  $\tilde{a}_k, \tilde{b}_k \in (2^{l+\lambda}, 2^{l+\lambda+1}) \cap \mathbb{Z}_N$  instead of  $a_k + r_k$  and  $b_k + s_k$  respectively, it would not be able to distinguish the difference. Similar arguments hold for the blinded messages  $c_i$  from the comparison protocol 7.

Hence, by processing multiple iterations, the target node can only obtain functions of the private data that involve at least one large enough random value, which does not break privacy.

We now build a simulator  $S_{\mathcal{T}}$  that applies the steps of the protocol on randomly generated values. As before, since Protocol 4 is secure in the semi-honest model (Proposition 2), we know that there exists a ppt simulator for the functionality of Protocol 4 on inputs  $I_{\mathcal{T}}^{-1} = \{sk_{\mathcal{T}}, sk_{DGK}\}$ . However, we need to show that we can simulate the functionality of consecutive calls of Protocol 4, or, equivalently, on one call of Protocol 4 but on inputs  $(I_{\mathcal{T}}^k, x^*)$ . Call such a simulator  $S_{\mathcal{T}}^k$ , that on the inputs  $(I_{\mathcal{T}}^k, x^*)$  should output a view that is statistically indistinguishable from  $V_{\mathcal{T}}^k(\bar{\mathcal{I}}^{k-1})$  in (30), for  $k = 0, \dots, K - 1$ . We already showed that although the target node has access to the output  $x^*$  and to blinded messages from all the iterations of Protocol 5, it cannot extract information from them or correlate the messages to the iteration they arise from, so  $x^*$  will only be relevant in the last simulator.

- Generate a  $\lambda + l$ -length random integer  $\tilde{\rho}_k$  and add  $2^l$  and obtain  $\tilde{z}_k$ ;
- Generate a random bit  $\tilde{t}_k$ ;
- Choose a random bit  $\delta_{\mathcal{T}}$ . If it is 0, then generate  $l$  non-zero values of  $2t$  bits, otherwise generate  $l - 1$  non-zero

random values and one 0 value. Those will be the  $\widetilde{m^{\text{comp}_k}}$  (see Protocol 7);

- Generate random integers of length  $l + \lambda + 1$   $\tilde{a}_k$  and  $\tilde{b}_k$ ;
- Compute  $\tilde{v}_k$  according to  $\tilde{t}_k$ ;
- Denote all Paillier and DGK generated coins by  $\widetilde{\text{coins}}$ ;
- Output  $\tilde{I}_{\mathcal{T}}^k := S_{\mathcal{T}}^k(I_{\mathcal{T}}^{k-1}, x^*) = (I_{\mathcal{T}}^{k-1}, \tilde{z}_k, \tilde{t}_k, \widetilde{m^{\text{comp}_k}}, \tilde{a}_k, \tilde{b}_k, \tilde{v}_k, \widetilde{\text{coins}})$ .

Finally, a trivial simulator  $\tilde{I}_{\mathcal{T}}^K := S_{\mathcal{T}}^K(I_{\mathcal{T}}^{K-1}, x^*)$  for  $V_{\mathcal{T}}^K(\bar{\mathcal{I}}^{K-1})$  is obtained by simply generating an encryption of  $x^*$  and outputting:  $(I_{\mathcal{T}}^{K-1}, \llbracket x^* \rrbracket)$ .

**Proposition C.2.**  $S_{\mathcal{T}}^k(I_{\mathcal{T}}^{k-1}, x^*) \stackrel{c}{\equiv} V^k(\bar{\mathcal{I}}^{k-1})$ , for  $k = 0, \dots, K$ .

*Proof.* For  $0 \leq k \leq K - 1$ ,  $(\text{coins}_{2k}, \text{coins}_{3k}) \stackrel{s}{\equiv} \widetilde{\text{coins}_k}$  because they are generated from the same distributions. Similarly,  $\tilde{z}_k \stackrel{s}{\equiv} z_k$  because  $\tilde{\rho}_k \stackrel{s}{\equiv} b_k - \frac{a_k}{\rho_k} + \rho_k$ . From the discussion above, the same holds for  $\widetilde{m^{\text{comp}_k}}$  and their counterparts  $m^{\text{comp}_k}$ . Furthermore,  $(\tilde{t}_k, \tilde{a}_k, \tilde{b}_k, \tilde{v}_k)$  are statistically indistinguishable from  $(t_k, \bar{a}_k, \bar{b}_k, v_k)$  due to the way they are generated, and  $\tilde{v}_k$  being consistent with  $\tilde{t}_k$ . Thus,  $S^k(I_{\mathcal{T}}^{k-1}, x^*) \stackrel{c}{\equiv} V^k(I_{\mathcal{T}}^{k-1} \cup \mathcal{I})$ .

For  $k = K$ ,  $S^K(I_{\mathcal{T}}^{K-1}, x^*) \stackrel{c}{\equiv} V^K(\bar{\mathcal{I}}^{K-1})$  trivially follows from the fact that both  $\llbracket x^* \rrbracket$  and  $\widetilde{\llbracket x^* \rrbracket}$  are decrypted in  $x^*$ .  $\square$

Hence, we obtained that  $\tilde{I}_{\mathcal{T}}^k \stackrel{c}{\equiv} I_{\mathcal{T}}^k$ , for  $k = 0, \dots, K$ . The essence of the proof of Proposition C.2 is that all the messages the target receives are blinded with large enough random noise and the comparison bits are randomized. This results in the messages being statistically indistinguishable from random values of the same length, which means that the extra information included in  $(I_{\mathcal{T}}^{k-1}, x^*)$  from the previous iterations cannot be used to extract other information about the current values at iteration  $k$ . The next corollary then follows from Proposition C.2:

**Corollary C.2.**  $S_{\mathcal{T}}^k(\tilde{I}_{\mathcal{T}}^{k-1}, x^*) \stackrel{c}{\equiv} S_{\mathcal{T}}^k(I_{\mathcal{T}}^{k-1}, x^*)$ , equivalent to  $S_{\mathcal{T}}^{k+1}(S_{\mathcal{T}}^k(I_{\mathcal{T}}^{k-1}, x^*), x^*) \stackrel{c}{\equiv} S_{\mathcal{T}}^{k+1}(V^k(\bar{\mathcal{I}}^{k-1}), x^*)$ , for any  $k = 1, \dots, K - 1$ .

Finally, we construct a simulator  $S_{\mathcal{T}}(I_{\mathcal{T}})$  for the execution of Protocol 5 and we show that its view will be statistically indistinguishable from  $V_{\mathcal{T}}(\mathcal{I})$ . To this end, we define the following sequence, from which we drop the input  $x^*$  to the simulators to not burden the notation:

$$\begin{aligned} V_{\mathcal{T}}(\mathcal{I}) &= H_0(\mathcal{I}) = V_{\mathcal{T}}^K(\bar{\mathcal{I}}^{K-1}) \\ H_1(\mathcal{I}, x^*) &= S_{\mathcal{T}}^K \circ V_{\mathcal{T}}^{K-1}(\bar{\mathcal{I}}^{K-2}) \\ H_2(\mathcal{I}, x^*) &= S_{\mathcal{T}}^K \circ S_{\mathcal{T}}^{K-1} \circ V_{\mathcal{T}}^{K-2}(\bar{\mathcal{I}}^{K-3}) \\ &\vdots \\ H_K(\mathcal{I}, x^*) &= S_{\mathcal{T}}^K \circ S_{\mathcal{T}}^{K-1} \circ \dots \circ S_{\mathcal{T}}^1 \circ V_{\mathcal{T}}^0(\mathcal{I}) \\ S_{\mathcal{T}}(I_{\mathcal{T}}) &= H_{K+1}(I_{\mathcal{T}}) = S_{\mathcal{T}}^K \circ S_{\mathcal{T}}^{K-1} \circ \dots \circ S_{\mathcal{T}}^1 \circ S_{\mathcal{T}}^0(I_{\mathcal{T}}) \end{aligned}$$

By transitivity,  $H_0$  and  $H_{K+1}$  are statistically indistinguishable if:

$$H_0 \stackrel{c}{\equiv} H_1 \stackrel{c}{\equiv} \dots \stackrel{c}{\equiv} H_k \stackrel{c}{\equiv} H_{k+1} \stackrel{c}{\equiv} H_{k+2} \stackrel{c}{\equiv} \dots \stackrel{c}{\equiv} H_{K+1}.$$

The result follows from induction on Corollary C.2. In conclusion, we obtain that  $S_{\mathcal{T}}(I_{\mathcal{T}}) \stackrel{c}{=} V_{\mathcal{T}}(\mathcal{I})$  which verifies that Protocol 5 achieves privacy with respect to the target node.

The proof of Theorem 1 is now complete.  $\square$

### B. Proof of Theorem 2

We will now show that any coalition consistent with the setup of Propositions 3, 4 and with the assumption that the cloud and target node do not collude does not gain any new information about the private data of the honest agents, other than what can be inferred solely from the inputs and outputs of the coalition. As mentioned in Section VI, the agents in a coalition only add their inputs to the view of the coalition, but do not have any messages in the protocol.

1) *Simulator for the cloud  $\mathcal{C}$  and  $\bar{p}$  agents  $\mathcal{A}_i$* : Consider the coalition between a set of agents  $\mathcal{A}_{i=1,\dots,\bar{p}}$  and the cloud  $\mathcal{C}$ , which has the inputs  $(\{b_i\}_{i=1,\dots,\bar{m}}, \{c_i\}_{i=1,\dots,\bar{n}}, A_{\mathcal{C}}, Q_{\mathcal{C}})$  and no output from the execution of Protocol 5. Since  $\bar{p} < p$ , the coalition is not able to compute  $\mu_1$  by itself, and the semantic security of the Paillier cryptosystem is again enough to ensure the privacy of the rest of the sensitive data. A simulator for this coalition can be constructed following the same steps in Section C-A2 on the augmented inputs defined above, from the fact that every value the cloud receives is encrypted by a semantically secure cryptosystem.

2) *Simulator for the target node  $\mathcal{T}$  and  $\bar{p}$  agents  $\mathcal{A}_i$* : Consider the coalition between a set of agents  $\mathcal{A}_{i=1,\dots,\bar{p}}$  and the target node  $\mathcal{T}$ , which has the inputs  $(\{b_i\}_{i=1,\dots,\bar{m}}, \{c_i\}_{i=1,\dots,\bar{n}}, sk_{\mathcal{T}}, sk_{DGG})$  and output  $(x^*)$  from the execution of Protocol 5. If both matrices  $Q_{\mathcal{C}}, A_{\mathcal{C}}$  are public, if there exists  $i \in \{1, \dots, \bar{p}\}$  such that  $a_i^T x^* < b_i$ , the coalition finds out  $\mu_i^* = 0$ , which comes from public knowledge in the KKT conditions (17). From this, the coalition is able to find some coins of the cloud:  $r_K$  and  $s_K$  associated to element  $i$ . However, these values are independent from the rest of the parameters and do not reveal any information about the private inputs of the parties. Apart from this, the coalition is not able to compute any private data of the honest parties from the execution of the protocol, due to the secure blinding. A simulator for this coalition can be build by following the same steps as described in Section C-A3 on the augmented inputs defined above.

The proof is now complete.  $\square$

### C. Sketch proof of Theorem 3

Following similar steps as in the proof of the main protocol in Appendix C-A, simulators for the involved parties can be constructed from their inputs and outputs. A simulator for the cloud randomly generates the messages and the indistinguishability between the views is guaranteed by the indistinguishability of the Paillier encryptions. The outputs of Protocol 6 to the target node are different than the original Protocol 4. This means that the outputs for the target node in Protocol 5 with the iterations as in Protocol 6 are  $(x^*, \{(r_k)_i; (\bar{\mu}_k)_i\}_{i=1,\dots,m, k=0,\dots,K-1})$ . In this case, a simulator for the target node simply outputs its inputs as messages in the view, and indistinguishability follows trivially. For the

multi-party privacy, simulators for the allowed coalitions can be easily constructed by expanding the simulators for the cloud and the target node with the auxiliary inputs from the agents in the coalition.  $\square$

### D. Proof of Proposition 3

The coalition has access to the following data, which is fixed:  $A_{\mathcal{C}}, Q_{\mathcal{C}}, x^*, \{b_i\}_{1,\dots,\bar{m}}, \{c_i\}_{i,\dots,\bar{n}}$ .

Proof of (1). We need to address two cases: the non-strict satisfaction of the constraints and the equality satisfaction of the constraints.

(I) Suppose there exists a solution  $\{b_i\}_{\bar{m}+1,\dots,m}$  and  $\{c_i\}_{\bar{n}+1,\dots,n}$  and  $\mu$  to the KKT conditions such that  $a_i^T x^* < b_i$  for some  $\bar{m} + 1 \leq i \leq m$ . In particular this implies that  $\mu_i = 0$ . Then define  $c' := c_{\mathcal{A}}$ ,  $\mu' := \mu$  and  $b'$  such that  $b'_j := b_j$  for all  $j \neq i$  and  $b'_i$  to take any value  $b'_i \geq a_i^T x^*$ . The new set of points  $(b', c', \mu')$  is also a solution to the KKT conditions, by construction.

(II) Alternatively, suppose there exists a solution  $\{b_i\}_{\bar{m}+1,\dots,m}$  and  $\{c_i\}_{\bar{n}+1,\dots,n}$  and  $\mu$  to the KKT conditions such that  $a_j^T x^* = b_j$  for all  $j = \bar{m} + 1, \dots, m$ . Consider there exists a vector  $\delta$  that satisfies  $\delta \succeq 0$  and  $A_{21}^T \delta = 0$ . Compute  $\epsilon \geq 0$  as:  $\epsilon = \min \left( \frac{\mu_k}{\delta_k} \right)_{\delta_k > 0, k = \bar{m} + 1, \dots, m}$ . Then, we construct

$\mu' := \mu - \epsilon \begin{bmatrix} 0 \\ \delta \end{bmatrix}$  that satisfies  $\mu' \succeq 0$  and  $\mu'_i = 0$  for some  $\bar{m} + 1 \leq i \leq m$  that is the argument of the above minimum.

Furthermore, define  $c' := c_{\mathcal{A}} + \epsilon \begin{bmatrix} 0 \\ A_{22}^T \delta \end{bmatrix}$  and  $b'$  such that  $b'_j := b_j$  for all  $j \neq i$  and  $b'_i$  to be any value  $b'_i > b_i$ . Then  $(b', c', \mu')$  is also a solution to the KKT conditions. More specifically, the complementarity slackness condition holds for all  $j = \bar{m} + 1, \dots, m$ :

$$\begin{aligned} \mu'_j (a_j^T x^* - b'_j) &= \mu'_j \underbrace{(a_j^T x^* - b_j)}_{=0} = 0, \quad j \neq i \\ \underbrace{\mu'_i}_{=0} (a_i^T x^* - b'_i) &= 0, \quad j = i. \end{aligned} \quad (32)$$

Then we can check the gradient condition:

$$\begin{aligned} Q_{\mathcal{C}} x^* + A_{\mathcal{C}}^T \mu' + c' &= \\ Q_{\mathcal{C}} x^* + A_{\mathcal{C}}^T \left( \mu - \epsilon \begin{bmatrix} 0 \\ \delta \end{bmatrix} \right) + \left( c + \epsilon \begin{bmatrix} 0 \\ A_{22}^T \delta \end{bmatrix} \right) & \quad (33) \\ = Q_{\mathcal{C}} x^* + A_{\mathcal{C}}^T \mu + c_{\mathcal{A}} &= 0. \end{aligned}$$

Hence,  $b' \neq b_{\mathcal{A}}$  satisfies the KKT conditions and the coalition cannot uniquely determine  $b_{\mathcal{A}}$ .

Proof of (2). Consider a solution  $\{b_i\}_{\bar{m}+1,\dots,m}$  and  $\{c_i\}_{\bar{n}+1,\dots,n}$  and  $\mu$  to the KKT conditions. For some  $\epsilon > 0$  define  $\mu' := \mu + \epsilon \begin{bmatrix} 0 \\ \delta \end{bmatrix}$  and  $c' := c_{\mathcal{A}} - \epsilon \begin{bmatrix} 0 \\ A_{22}^T \delta \end{bmatrix}$ . Define  $b'$  such that for all  $j$ , it holds that  $b'_j = a_j^T x^*$ . Then  $(b', c', \mu')$  is also a solution to the KKT conditions. Specifically, it follows that  $\mu' \geq 0$ . Moreover the complementarity slackness condition holds by construction of  $b'$ , and as before the gradient condition holds. Hence,  $c' \neq c_{\mathcal{A}}$  satisfies the KKT solution, and the coalition cannot uniquely determine  $c_{\mathcal{A}}$ .  $\square$