2020 American Control Conference
Denver, CO, USA, July 1-3, 2020

Privacy-Preserving Transactive Energy System

Yang Lu, Jianming Lian and Minghui Zhu

Abstract— **In this paper, the privacy issue of the recently proposed transactive energy system for electric power system is investigated for the first time. It is identified that the private information of individual market participants will be subject to the risk of leakage during the market interactions. In order to enable the feature of privacy preservance for market participants, a homomorphic encryption-based approach is developed to augment the existing design of transactive energy system. The proposed privacy-preserving design based on the Paillier encryption scheme is then demonstrated on a transactive energy system that coordinates and controls residential air conditioners under the same feeder to manage the feeder congestion. The simulation results confirm the effectiveness of the proposed design in protecting the privacy of individual market participants without affecting the overall system performance.**

I. Introduction

Over the past decade, a new class of engineering systems that is referred to as the cyber-physical system (CPS) has rapidly emerged. Typical examples of CPS include smart buildings, smart grids, process control systems, intelligent transportation systems, autonomous vehicles, robotic systems and so on. For any CPS, a large number of geographically dispersed entities are coordinated to achieve network-wide objectives. Distributed data sharing, which is necessary to enable the required coordination, actually raises the significant concern that private information of legitimate entities could be leaked to adversarial ones.

In the literature, there have been many techniques proposed to protect the CPS privacy [1]. Mutual information has been used to define data privacy in smart grids [2]. Given specific statistical models of source data and auxiliary information, the posterior information entropy of private data is quantified and minimized so as to minimize information leakage. Obfuscation has been used to protect coefficient privacy in cloud computing in centralized optimal power flow problems [3]. The legitimate problem holder obfuscates the optimization problem by an invertible transformation and sends the obfuscated problem to the cloud. The cloud solves the obfuscated problem and sends the solution back to the legitimate problem holder, who can retrieve an optimal solution of the original optimization problem by inverting the obfuscation transformation. Differential privacy has been applied to distributed optimization [4]–[6], consensus [7], and filtering [8]. Differentially private schemes add random noises into individuals' data such that an adversary with arbitrary auxiliary information cannot infer an individual's participation. However, there is a fundamental utility-privacy trade-off for differentially private schemes due to the introduction of random noises [9].

For electric power system, transactive control is emerging as a new type of approaches for integrating many distributed energy resources into grid operations. It engages individual resources as market participants (energy suppliers or customers) through market interactions, and uses the market clearing prices to achieve their optimal coordination for both local and global objectives. Various transactive energy systems (TESs) have been proposed (see, for example, [10]–[12]). TES is actually a typical CPS, where the cyber space determines market clearing prices and the physical world performs control tasks. Hence, TESs share the typical privacy issues of general CPSs. Based on the way of information exchange, there are two types of market clearing: hierarchical and distributed. For the hierarchical clearing, also known as auction-based clearing, each resource submits to the coordinator its entire supply or demand curve (the relation between energy price and supply or demand). The coordinator can then use the received curves to infer individual resources' private information, e.g., business secrets and personal preferences. The distributed clearing is more privacy-aware as each resource only reports those points of its supply or demand curve required by the coordinator, rather than the entire curve. However, the coordinator could intentionally require a large number of points such that it can essentially recover the entire curve. The above privacy issue necessitates novel TESs that can realize market-based coordination and simultaneously protect privacy of individual agents.

In this paper, homomorphic encryption (HE) is applied to address the privacy issue of TESs. HE is a cryptographic technique allowing certain algebraic operations to be carried out on ciphertexts, thus generating an encrypted result which, when decrypted, matches the result of operations performed on plaintexts. A significant advantage of HE is that it can achieve perfect correctness in secure multiparty computation, i.e., after the multiparty computation, each party can obtain the correct result of its target computation, and meanwhile no information of its private data is disclosed to any other party. HE has been increasingly adopted to ensure data privacy in multiparty computation in various control and CPS problems, e.g., potential games [13], distributed optimization [14], quadratic programs [15], and consensus [16]. Because TESs

This work was supported by the Transactive Control Program at the Pacific Northwest National Laboratory (PNNL) funded by the U.S. Department of Energy. PNNL is operated for the U.S. Department of Energy by Battelle Memorial Institute under Contract DE-AC05-76RL01830. M. Zhu was also partially supported by the NSF CAREER award ECCS-1846706.

Y. Lu and M. Zhu are with the School of Electrical Engineering and Computer Science, Pennsylvania State University, University Park, PA 16802, USA (Email: {yml5046, muz16}@psu.edu).

J. Lian is with the Optimization and Control Group, Pacific Northwest National Laboratory, Richland, WA 99354, USA (*Corresponding author*; Email: jianming.lian@pnnl.gov).

involve multiparty computation related to the coordination, HE is a good candidate to address its privacy issue.

To cope with the computation (addition) involved in market-based coordination, this paper adopts the Paillier encryption scheme [17], which is an additive HE scheme. Each agent encrypts each sampling point of its supply or demand curve using the coordinator's public key and then sends the encrypted data to a third party. The third party performs certain aggregation operations over the encrypted data and sends the aggregated encrypted data to the coordinator, who then uses its private key to perform decryption. By the homomorphic property of the Paillier encryption scheme, the decrypted data is just the original aggregated sampling point, by which the coordinator can determine the correct market clearing price. In the above process, the coordinator only learns the aggregated sampling point, but does not know individual sampling points. Moreover, without knowing the coordinator's private key, any other entity cannot learn anything about an agent's individual sampling points by observing the ciphertext. The effectiveness of the proposed design is verified via simulation on a residential air conditioners (ACs) coordination problem.

Notations. Denote by $\mathbb{R}$ and $\mathbb{N}$ the sets of real and natural numbers, respectively. Given a positive integer n, let $\mathbb{Z}_n \triangleq \{0, 1, \cdots, n-1\}$ and let $\mathbb{Z}_n^*$ be the set of positive integers that are smaller than and co-prime to n. Given positive integers x and y, denote by $\gcd(x, y)$ and $\mathrm{lcm}(x, y)$ the greatest common divisor and the least common multiple of x and y, respectively. Given a set S, denote by $|S|$ its cardinality.

II. Preliminaries on Encryption

In this section, we first introduce the structure of public-key encryption and some standard privacy definitions, and then illustrate the Paillier encryption scheme. More detailed discussions of this section can be found in [17], [18].

A public-key encryption scheme is a triple, (G, E, D), of probabilistic polynomial-time (PPT) algorithms, such that: (1) on an input of security parameter n (the key length), algorithm G (key-generator) outputs a pair of keys, and (2) for every pair (e, d) in the range of $G(n)$ and every admissible message m, algorithms E (encryption) and D (decryption) satisfy $\Pr[D(d, E(e, m)) = m] = 1$. In the above, e is called the public key and d is called the private key. In the following, we denote the first and second elements of $G(n)$ by $G_1(n)$ and $G_2(n)$, respectively, i.e., $G(n) = (G_1(n), G_2(n))$. In order to define the privacy of a public-key encryption scheme, the notion of computational indistinguishability is first introduced.

Definition 2.1 ([18]): Let $X = \{X_n\}_{n\in\mathbb{N}}$ and $Y = \{Y_n\}_{n\in\mathbb{N}}$ be two probability ensembles, where X_n and Y_n are random variables for each $n \in \mathbb{N}$. We say that X and Y are computationally indistinguishable, denoted by $X \stackrel{c}{\equiv} Y$, if for every non-uniform PPT distinguisher D, every positive polynomial p, and all sufficiently large n, it holds that $|\Pr[D(X_n) = 1] - \Pr[D(Y_n) = 1]| < 1/p(n)$.

The privacy of a public-key encryption scheme is defined by the following notion of semantic security.

Algorithm 1: Key generation algorithm

Syntax: $(\alpha, \beta, \nu, \pi) = \mathrm{Alg}_{\mathrm{key}}(n)$.

The executor randomly chooses prime numbers p and q such that $\gcd(pq, (p-1)(q-1)) = 1$ and $|\alpha| = n$ with $\alpha = pq$; computes $\nu = \mathrm{lcm}(p-1, q-1)$; randomly selects an integer $\beta \in \mathbb{Z}^*_{\alpha^2}$ such that the following modular multiplicative inverse π exists: $\pi = (\frac{(\beta^\nu \mod \alpha^2)-1}{\alpha})^{-1} \mod \alpha$.

Definition 2.2 ([18]): A public-key encryption scheme (G, E, D) is semantically secure if for every e in the range of G_1, every admissible plaintexts x, y, it holds that $\{e, E(e, x)\} \stackrel{c}{\equiv} \{e, E(e, y)\}$.

Definition 2.2 states that, even with the public key, it is infeasible to distinguish the encryptions of any two plaintexts. This definition is dedicated to public-key encryption schemes. We still need a privacy notion for general problem settings to enable privacy analysis when applying public-key encryption schemes to different applications. Such a privacy notion is defined for the general setting of secure multiparty computation (SMC), which is illustrated in the following.

Consider a set of M parties. Each party i has a private input x_i and aims to compute a function $f_i(x)$ with $x = (x_1, \cdots, x_M)$. If a party does not have a private input or a function to compute, then we say that its input or output is null. Let Π be an algorithm to compute $f = (f_1, \cdots, f_M)$. The *view* of a party i during an execution of Π on a joint input x, denoted by $\mathrm{VIEW}_i^\Pi(x)$, is defined as $\mathrm{VIEW}_i^\Pi(x) \triangleq \{x_i, r_i, m_1^i, \cdots, m_{t_i}^i\}$, where r_i is the vector consisting of its internal random choices, t_i is the number of messages it receives during the execution of Π, and m_ℓ^i is the ℓ-th message it receives. The privacy of an algorithm in the general setting of SMC is defined next.

Definition 2.3 ([18]): Let Π be an algorithm for computing $f(x)$. We say that Π is privacy-preserving if there exists a PPT algorithm S such that for each $i \in \{1, \cdots, M\}$ and every admissible x, it holds that $S(i, x_i, f_i(x)) \stackrel{c}{\equiv} \mathrm{VIEW}_i^\Pi(x)$, where $S(i, \cdot)$ denotes the set of messages party i can see after the execution of S.

Definition 2.3 states that each party i's view during the execution of Π can be efficiently simulated by only using its own input x_i and output $f_i(x)$.

Paillier encryption. The Paillier encryption scheme is a public-key encryption scheme which, in addition, has an additive homomorphic property. It consists of key generation, encryption and decryption operations, detailed next.

- Key generation: A set of keys $(\alpha, \beta, \nu, \pi)$ is generated by Algorithm 1, where n is the parameter to set the key length, (α, β) are public keys, and (ν, π) are private keys.
- Encryption: A plaintext $pt \in \mathbb{Z}_\alpha$ is encrypted as ct with the public keys (α, β) by Algorithm 2.
- Decryption: A ciphertext $ct \in \mathbb{Z}_{\alpha^2}$ is decrypted as pt with the private keys (ν, π) by Algorithm 3.

The correctness, privacy and homomorphic property of the Paillier encryption scheme are summarized as follows:

Algorithm 2: Encryption algorithm

Syntax: $ct = \text{Alg}_{\text{enc}}(\alpha, \beta, pt)$.

The executor selects a random integer $r \in \mathbb{Z}_\alpha^*$ and computes $ct = \beta^{pt} \cdot r^\alpha \mod \alpha^2$.

Algorithm 3: Decryption algorithm

Syntax: $pt = \text{Alg}_{\text{dec}}(\alpha, \nu, \pi, ct)$.

The executor computes $pt = \frac{(ct^\nu \bmod \alpha^2)-1}{\alpha}\pi \bmod \alpha$.

(1) Correctness: $\text{Alg}_{\text{dec}}(\alpha, \nu, \pi, \text{Alg}_{\text{enc}}(\alpha, \beta, pt)) = pt$.

(2) Privacy: If the decisional composite residuosity assumption (DCRA)[1] holds, then the Paillier encryption scheme is semantically secure.

(3) Homomorphic property: Given any $pt_1, \cdots, pt_m \in \mathbb{Z}_\alpha$. If $\sum_{\ell=1}^m pt_\ell \in \mathbb{Z}_\alpha$, then $\text{Alg}_{\text{dec}}\left(\alpha, \nu, \pi, \prod_{\ell=1}^m \text{Alg}_{\text{enc}}(\alpha, \beta, pt_\ell)\right) = \sum_{\ell=1}^m pt_\ell$.

III. Transactive Energy System

In this section, the basic concept of TES will be briefly reviewed. Two widely used market clearing approaches adopted by the TES will also be illustrated.

The TES can be modeled as a multi-agent system, where a group of resources are coordinated and controlled to achieve certain social objectives while respecting local preferences and constraints. Three different types of agents including coordinator (CO), supplier, and customer are involved in the TES. The coordinator represents the market operator, the supplier is the electricity seller, and the customer is the electricity buyer. The underlying coordination and control usually have a hierarchical structure as shown by Fig. 1. The TES is usually applied to solve the resource allocation problem, where the market operator achieves the optimal resource allocation by properly setting the resource price, referred to as the market clearing price. In the following, the mathematical formulation of a typical TES in electric power systems is given for the illustration purpose.

Denote by $\mathcal{V}_s$ and $\mathcal{V}_d$ the set of suppliers and the set of customers, respectively. The subscripts "s" and "d" indicate "supply" and "demand", respectively. Let $N_s = |\mathcal{V}_s|$, $N_d = |\mathcal{V}_d|$, and $\mathcal{V} \triangleq \mathcal{V}_s \cup \mathcal{V}_d$. Each supplier $i \in \mathcal{V}_s$ aims to find an optimal supply that maximizes its profit, defined as the earnings in energy selling minus the costs in energy purchasing. The profit optimization problem of supplier $i \in \mathcal{V}_s$ is formulated as $\max_{p_i^s \in \mathcal{L}_i^s} \lambda p_i^s - C_i(p_i^s)$, where p_i^s is its supply, $C_i : \mathbb{R} \to \mathbb{R}$ is its cost function, λ is the energy price, and $\mathcal{L}_i^s$ is the feasible set of p_i^s. Each customer $i \in \mathcal{V}_d$ aims to find an optimal demand that maximizes its utility, defined as the benefit in energy usage minus the energy cost. The utility optimization problem of customer $i \in \mathcal{V}_d$ is formulated as $\max_{p_i^d \in \mathcal{L}_i^d} U_i(p_i^d) - \lambda p_i^d$, where p_i^d is its demand, $U_i : \mathbb{R} \to \mathbb{R}$ is its benefit function, and $\mathcal{L}_i^d$ is the feasible set of p_i^d.

[1]DCRA: Given a composite C and an integer z, it is computationally intractable to decide whether z is a C-residue modulo C^2 or not, i.e., whether there exists y such that $z = y^C \mod C^2$. It is widely believed that the DCRA is true.

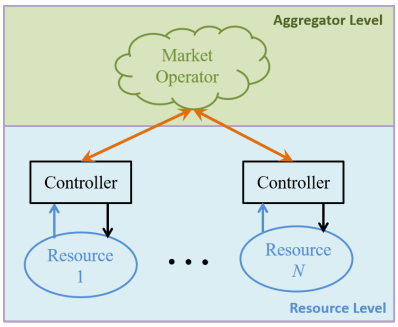

Fig. 1. Hierarchical structure of TES for coordination and control.

The coordinator aims to obtain an aggregated optimal allocation that maximizes the social welfare. The bi-level optimization problem of the coordinator is formulated as

$$\begin{aligned} &\max_{\lambda \in \mathbb{R}} \sum_{i \in \mathcal{V}_d} U_i(p_i^{d*}(\lambda)) - \sum_{i \in \mathcal{V}_s} C_i(p_i^{s*}(\lambda)) \\ &\text{s.t.} \quad p_i^{s*}(\lambda) = \operatorname{argmax}_{p_i^s \in \mathcal{L}_i^s} \lambda p_i^s - C_i(p_i^s),\ \forall i \in \mathcal{V}_s, \\ &\qquad p_i^{d*}(\lambda) = \operatorname{argmax}_{p_i^d \in \mathcal{L}_i^d} U_i(p_i^d) - \lambda p_i^d,\ \forall i \in \mathcal{V}_d, \\ &\qquad \sum_{i \in \mathcal{V}_s} p_i^{s*}(\lambda) = \sum_{i \in \mathcal{V}_d} p_i^{d*}(\lambda). \end{aligned} \tag{1}$$

In problem (1), the function $p_i^{s*}(\cdot)$ (resp. $p_i^{d*}(\cdot)$) is called the supply (resp. demand) curve. Both hierarchical and distributed market clearing approaches have been widely used to determine the optimal solution λ^* of problem (1).

Hierarchical market clearing. Hierarchical clearing is also referred to as the auction-based clearing. Each agent $i \in \mathcal{V}$ submits its entire supply or demand curve to the coordinator, who determines the clearing price as λ^* such that $\sum_{i \in \mathcal{V}_s} p_i^{s*}(\lambda^*) = \sum_{i \in \mathcal{V}_d} p_i^{d*}(\lambda^*)$.

Distributed market clearing. Distributed clearing works in an iterative manner. At each iteration k, the coordinator broadcasts the current clearing price estimate $\lambda(k)$ to all the agents. Each supplier $i \in \mathcal{V}_s$ (resp. customer $i \in \mathcal{V}_d$) reports $p_i^s(k) = p_i^{s*}(\lambda(k))$ (resp. $p_i^d(k) = p_i^{d*}(\lambda(k))$) to the coordinator, who then updates the clearing price estimate for the next iteration until convergence is reached.

IV. Problem Statement

In this section, we identify the privacy issue of TES and clarify our objective.

In the hierarchical approach, the agents submit their entire supply or demand curves to the coordinator. With this information, the coordinator or an eavesdropper can infer individual cost or benefit functions. In fact, the inverse supply or demand function is just the derivative of the corresponding cost or benefit function [19]. Hence, individual cost or benefit functions can be recovered by integrating the inverse of the corresponding supply or demand functions. This could expose much about individuals' business secrets (for suppliers) or personal preferences (for customers).

The distributed approach can to some extent mitigate the privacy issue as the agents do not submit to the coordinator their entire supply or demand curves, but only those required

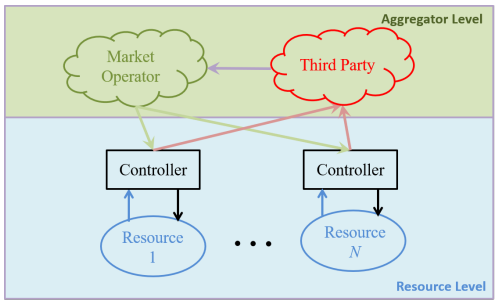

Fig. 2. Proposed practical deployment for TES.

points. However, the coordinator could make use of the iterative nature of the distributed approach to intentionally broadcast a large number of prices covering the whole admissible range and obtains the agents' responses. In this way, the coordinator or an eavesdropper could recover the agents' supply or demand curves arbitrarily well.

Objective. We aim to develop a privacy-preserving algorithm that simultaneously satisfies the following properties:

(1) Correctness: The coordinator can determine the correct clearing price λ^* such that $\sum_{i\in\mathcal{V}_s} p_i^{s*}(\lambda^*) = \sum_{i\in\mathcal{V}_d} p_i^{d*}(\lambda^*)$;

(2) Privacy preservation: After the execution of the algorithm, for each supplier $i \in \mathcal{V}_s$ (resp. customer $i \in \mathcal{V}_d$), no other entity can infer $p_i^{s*}(\lambda)$ (resp. $p_i^{d*}(\lambda)$) for any λ.

In the next section, we will propose a practical deployment for TES under which HE can be applied to simultaneously achieve the above two objectives. For the purpose of illustration, we will only consider the hierarchical clearing. However, it is straightforward to extend the proposed design to the distributed clearing.

V. Priacy-preserving design

In this section, an HE-based privacy-preserving TES design is developed. We first propose a practical deployment of TES to enable the usage of HE and illustrate the attacker model. After that, we present the privacy-preserving design for the auction-based approach.

A. Practical Deployment

The privacy concern requires that the coordinator should obtain aggregated curves without knowing individual ones. Homomorpihc encryption is a promising technique to fulfill this requirement. For the privacy concern, this technique requires that the entity who receives individual ciphertexts and carries out algebraic operations to be different from the entity who performs decryptions. Hence, to enable the usage of HE, we introduce an additional independent third party (TP) as the entity who receives individual ciphertexts from market participants and performs encrypted aggregations for the coordinator. The proposed practical deployment is shown in Fig. 2, in which we assume that there is a communication link (i, TP) between each agent $i \in \mathcal{V}$ and the third party, and a communication link $(\mathrm{TP}, \mathrm{CO})$ between the third party and the coordinator.

Attacker model. We assume that any market participant $i \in \mathcal{V}\cup\{\mathrm{CO}, \mathrm{TP}\}$ is semi-honest, i.e., it correctly follows the designed algorithm but attempts to use its received messages throughout the execution of the algorithm to infer other participants' private data ([20], pp-20). The communication links are non-confidential and could be eavesdropped. An eavesdropper could be the third party, any supplier or customer, or an extraneous entity. We assume that the coordinator is not an eavesdropper and does not collaborate with any other entity.

B. Algorithm design

This subsection presents the proposed privacy-preserving auction-based algorithm. Denote by $\lambda_{\min}$ and $\lambda_{\max}$ the lower and upper bounds of energy price, respectively. Denote by τ the sampling length and N_p the number of sampling points. For each supplier $i \in \mathcal{V}_s$ (resp. customer $i \in \mathcal{V}_d$), denote by $p_{i\ell}^{s*}$ (resp. $p_{i\ell}^{d*}$) its ℓ-th sampling point, i.e., $p_{i\ell}^{s*} = p_i^{s*}(\lambda_{\min} + \ell\tau)$ (resp. $p_{i\ell}^{d*} = p_i^{d*}(\lambda_{\min} + \ell\tau)$). Denote by $\sigma \in \mathbb{N}$ the precision level of the sampling points, that is, for any $p_{i\ell}^{s*}$ and $p_{i\ell}^{d*}$, only σ decimal fraction digits are kept. Assume that the coordinator and all the suppliers (resp. customers) know a strict upper bound δ_s (resp. δ_d) of individual supply (resp. demand), i.e., $\delta_s > p_i^s$ for all $i \in \mathcal{V}_s$ and all $p_i^s \in \mathcal{L}_i^s$ (resp. $\delta_d > p_i^d$ for all $i \in \mathcal{V}_d$ and all $p_i^d \in \mathcal{L}_i^d$).

Our privacy-preserving auction-based algorithm, Algorithm 4, is based on the Paillier encryption scheme (please refer to Section II), and informally stated next.

Algorithm 4: Privacy-preserving auction

1 **Key generation**
The CO runs $(\alpha, \beta, \nu, \pi) = \mathrm{Alg}_{\mathrm{key}}(n)$ such that $\alpha > \max\{10^\sigma N_s\delta_s, 10^\sigma N_d\delta_d\}$, broadcasts (α, β) and keeps (ν, π) private to itself;
for $\ell = 1$; $\ell \le N_p$; $\ell = \ell + 1$ **do**
2 **Encryption**
Each supplier $i \in \mathcal{V}_s$ runs $y_{i\ell}^s = \mathrm{Alg}_{\mathrm{enc}}$ $(\alpha, \beta, 10^\sigma p_{i\ell}^{s*})$ and sends $y_{i\ell}^s$ to the TP;
Each customer $i \in \mathcal{V}_d$ runs $y_{i\ell}^d = \mathrm{Alg}_{\mathrm{enc}}$ $(\alpha, \beta, 10^\sigma p_{i\ell}^{d*})$ and sends $y_{i\ell}^d$ to the TP;
3 **Computation over ciphertexts**
TP computes $y_\ell^s = \prod_{i\in\mathcal{V}_s} y_{i\ell}^s \mod \alpha^2$ and $y_\ell^d = \prod_{i\in\mathcal{V}_d} y_{i\ell}^d \mod \alpha^2$, sends (y_ℓ^s, y_ℓ^d) to CO;
4 **Decryption**
CO runs $\hat{y}_\ell^s = \mathrm{Alg}_{\mathrm{dec}}(\alpha, \nu, \pi, y_\ell^s)\ /10^\sigma$ and $\hat{y}_\ell^d = \mathrm{Alg}_{\mathrm{dec}}(\alpha, \nu, \pi, y_\ell^d)/10^\sigma$;
5 **Setting clearing price**
CO sets $\lambda^* = \lambda_{\min} + \ell\tau$ such that $\hat{y}_\ell^s = \hat{y}_\ell^d$, and sends λ^* to each agent $i \in \mathcal{V}$.

At step 1, the coordinator generates a set of keys by the Paillier key-generation algorithm. The public keys are broadcasted while the private keys are kept private to itself. The bound on α is to guarantee decryption correctness. Roughly speaking, to ensure decryption correctness, α must be larger than the computing result. At step 2, each supplier

$i \in \mathcal{V}_s$ (resp. customer $i \in \mathcal{V}_d$) encrypts its sampling point $10^\sigma p_{i\ell}^{s*}$ (resp. $10^\sigma p_{i\ell}^{d*}$) by the Paillier encryption algorithm with the public keys (α, β), and sends the ciphertext $y_{i\ell}^s$ (resp. $y_{i\ell}^d$) to the third party. Notice that $10^\sigma p_{i\ell}^{s*}$ and $10^\sigma p_{i\ell}^{d*}$ are non-negative integers. At step 3, the third party performs computations over received ciphertexts according to the homomorphic property of the Paillier encryption scheme, that is, multiplication of ciphertexts provides an encryption of sum of plaintexts. Hence, y_ℓ^s and y_ℓ^d are encryptions of the ℓ-th sampling points of the aggregated supply and demand curves, respectively. The third party then sends y_ℓ^s and y_ℓ^d to the coordinator. At step 4, the coordinator decrypts y_ℓ^s and y_ℓ^d by the Paillier decryption algorithm with its public key α and private keys (ν, π), and transforms the decrypted results back to real numbers via dividing them by 10^σ. At step 5, the coordinator sets and broadcasts the clearing price λ^*.

Algorithm 4 has the following properties:

(1) Correctness: For each $\ell \in \{1, \cdots, N_p\}$, it holds that $\hat{y}_\ell^s = \sum_{i\in\mathcal{V}_s} p_i^{s*}(\lambda_{\min}+\ell\tau)$ and $\hat{y}_\ell^d = \sum_{i\in\mathcal{V}_d} p_i^{d*}(\lambda_{\min}+\ell\tau)$.

The correctness property states that $\hat{y}_\ell^s$ and $\hat{y}_\ell^d$ are just the ℓ-th sampling points of the original aggregated supply and demand curves, respectively. This property directly follows the homomorphic property of the Paillier encryption scheme. Since λ^* is set as $\lambda^* = \lambda_{\min}+\ell\tau$ such that $\hat{y}_\ell^s = \hat{y}_\ell^d$, the correctness property leads to $\sum_{i\in\mathcal{V}_s} p_i^{s*}(\lambda^*) = \sum_{i\in\mathcal{V}_d} p_i^{d*}(\lambda^*)$. Hence, optimal market-based coordination is achieved.

(2) Privacy preservation: If the DCRA holds, then Algorithm 4 is privacy-preserving in the sense of Definition 2.3.

The privacy preservation property directly follows the semantic security of the Paillier encryption scheme. In our problem setting, since the clearing price λ^* is public information, it is included in every entity's output. For each supplier $i \in \mathcal{V}_s$, its input is $(p_{i\ell}^{s*})_{\ell\in\{1,\cdots,N_p\}}$ and output is λ^*. For each customer $i \in \mathcal{V}_d$, its input is $(p_{i\ell}^{d*})_{\ell\in\{1,\cdots,N_p\}}$ and output is λ^*. The coordinator's input is null and output is $((\sum_{i\in\mathcal{V}_s} p_{i\ell}^{s*}, \sum_{i\in\mathcal{V}_d} p_{i\ell}^{d*})_{\ell\in\{1,\cdots,N_p\}}, \lambda^*)$. For the third party or an extraneous eavesdropper, its input is null and output is λ^*. By Definition 2.3, after the execution of Algorithm 4, each entity only knows its own input and output. Hence, each agent $i \in \mathcal{V}$ only knows its own supply or demand curve and the market clearing price; the coordinator only knows the aggregated supply and demand curves and the market clearing price; the third party or an extraneous eavesdropper only knows the market clearing price. Therefore, any agent's individual supply or demand curve is not known to any other entity and privacy preservation is achieved.

VI. Case Studies

In this section, the proposed privacy-preserving design is tested on a TES that coordinates and controls residential air conditioners to manage the feeder congestion.

We consider the real-time electricity allocation of a distribution feeder on a hot summer day (August 16, 2009) for Columbus, Ohio, USA. The weather data and the Typical Meteorological Year (TMY2) data are adopted from [21] and [22]. The wholesale energy price is adopted from the PJM market [23] and it is modified to a retail rate in \$/kWh plus a retail modifier as defined by American Electric Power (AEP)'s tariff [24]. We define this retail price as the base price. The distribution feeder capacity limit is 3.5 MW. There are 1000 residential ACs under the feeder. The feeder is both the coordinator and the (only) supplier, and each residential AC is a customer. In each market cycle, the feeder aims to obtain the aggregated demand curve and compares it with the feeder capacity limit to determine the market clearing price. If there is no congestion, then the clearing price is set to the base price. If there is congestion, the clearing price is set as the price corresponding to the feeder capacity limit on the aggregated demand curve. In the simulation, the price range is between $\lambda_{\min} = \$0$ to $\lambda_{\max} = \$1$ and the sampling length is $\tau = \$0.01$. The length of a market cycle is 5 minutes and there is 288 market cycles in total for one day.

We simulate the above problem for a whole day. A second-order equivalent thermal parameter (ETP) model is used to capture the load dynamics of the ACs. Deteailed description of the ETP model parameters can be found in [25]. Fig 3 shows the evolution of feeder power within 24 hours. The trajectory of feeder power with control (the solid blue line) is derived under the proposed privacy-preserving algorithm. Fig 3 verifies that our algorithm maintains optimal market-based coordination. Fig. 4 shows the aggregated demand curve at the 220-th market cycle. Denote by $p^{d*}(\lambda)$ the aggregated demand function, i.e., $p^{d*}(\lambda) \triangleq \sum_{i\in\mathcal{V}_d} p_i^{d*}(\lambda)$. We simulate the auction-based scheme both with and without our privacy-preserving design and denote the aggregated demand functions derived in the two cases by $p_{\text{privacy}}^{d*}(\lambda)$ and $p_{\text{plain}}^{d*}(\lambda)$, respectively. In Fig. 4, the image of $p_{\text{privacy}}^{d*}(\lambda)$ (the solid blue line) shows the shape of the aggregated demand curve, and the image of $|p_{\text{privacy}}^{d*}(\lambda) - p_{\text{plain}}^{d*}(\lambda)|$ (the dot red line), which is constant at 0, shows that $p_{\text{privacy}}^{d*}(\lambda)$ is exactly equal to $p_{\text{plain}}^{d*}(\lambda)$ at all values of λ, which verifies the correctness of Algorithm 4. In Fig. 5, the left subfigure shows agent 100's demand curve at the 200-th market cycle, and the right subfigure shows its encryption under 500 bits of key length. Fig. 5 visually illustrates the privacy preservation of Algorithm 4, as the points of the encrypted demand curve look like pure random numbers within a large interval. Table I lists the running time under different key lengths. The time in the second column is the average time per agent per market cycle, and the time in the second and third columns is the average time per market cycle. We can see that even for the case where the key length has 4000 bits, each participant's running time per market cycle is still much less than 5 minutes. This verifies that our privacy-preserving algorithm is efficient enough for TESs.

VII. Conclusion

This paper studies privacy issue of TESs. An HE-based algorithm is developed to simultaneously achieve optimal market-based coordination and privacy preservation. The effectiveness of the proposed algorithm is verified by a residential ACs coordination problem.

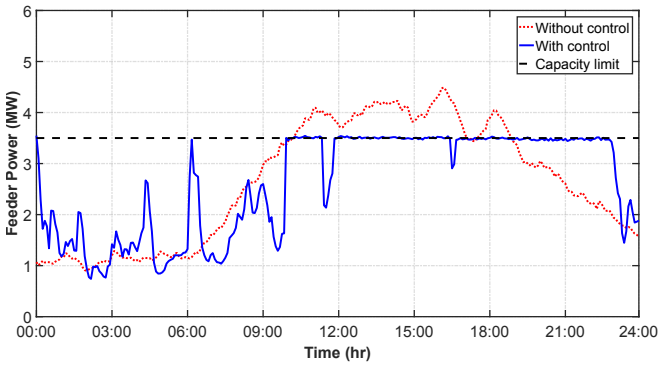

Fig. 3. Evolution of feeder power within 24 hours

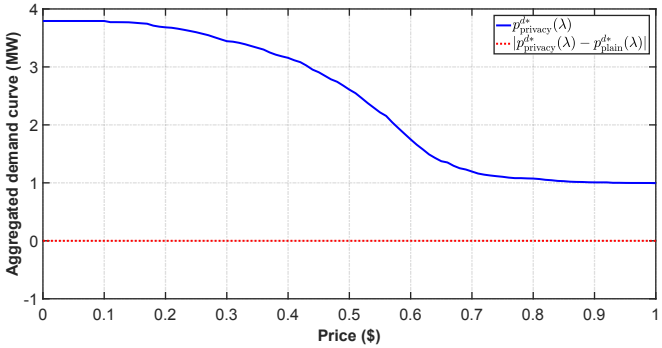

Fig. 4. Aggregated demand curve at market cycle 220

TABLE I

COMPUTATIONAL EFFICIENCY

Key length (bit)	Agent (s)	TP (s)	CO (s)
500	0.57	11.99	0.54
1000	2.52	16.50	3.29
1500	7.56	26.44	9.85
2000	16.78	42.06	21.92
2500	32.45	69.59	42.42
3000	55.04	106.43	75.82
3500	86.12	157.02	112.95
4000	127.76	221.99	169.42

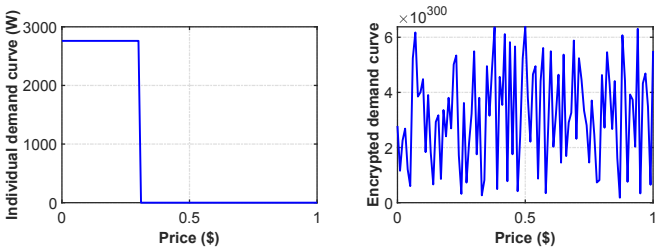

Fig. 5. Agent 100's demand curve at market cycle 200

REFERENCES

[1] Y. Lu and M. Zhu, "A control-theoretic perspective on cyber-physical privacy: Where data privacy meets dynamic systems," *Annual Reviews in Control*, vol. 47, pp. 423–440, 2019.

[2] O. Tan, D. Gunduz, and H. V. Poor, "Increasing smart meter privacy through energy harvesting and storage devices," *IEEE Journal on Selected Areas in Communications*, vol. 31, no. 7, pp. 1331–1341, 2013.

[3] A. R. Borden, D. K. Molzahn, B. C. Lesieutre, and P. Ramanathan, "Power system structure and confidentiality preserving transformation of optimal power flow problem," in *Fifty-first Annual Allerton Conference*, 2013, pp. 1021–1028.

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

[5] 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.

[6] M. T. Hale and M. Egerstedt, "Cloud-enabled differentially private multiagent optimization with constraints," *IEEE Transactions on Control of Network Systems*, vol. 5, no. 4, pp. 1693–1706, 2018.

[7] Z. Huang, S. Mitra, and G. Dullerud, "Differentially private iterative synchronous consensus," in *Proceedings of the 2012 ACM workshop on privacy in the electronic society*, 2012, pp. 81–90.

[8] J. L. Ny and G. J. Pappas, "Differentially private filtering," *IEEE Transactions on Automatic Control*, vol. 59, no. 2, pp. 341–354, February 2014.

[9] Q. Geng and P. Viswanath, "Optimal noise adding mechanisms for approximate differential privacy," *IEEE Transactions on Information Theory*, vol. 62, no. 2, pp. 952–969, 2016.

[10] J. de Hoog, T. Alpcan, M. Brazil, D. A. Thomas, and I. Mareels, "A market mechanism for electric vehicle charging under network constraints," *IEEE Transactions on Smart Grid*, vol. 7, no. 2, pp. 827–836, March 2016.

[11] Y. Xu, N. Li, and S. H. Low, "Demand response with capacity constrained supply function bidding," *IEEE Transactions on Power Systems*, vol. 31, no. 2, pp. 1377–1394, March 2016.

[12] K. Dehghanpour and H. Nehrir, "An agent-based hierarchical bargaining framework for power management of multiple cooperative microgrids," *IEEE Transactions on Smart Grid*, vol. 10, no. 1, pp. 514–522, January 2019.

[13] Y. Lu and M. Zhu, "Secure cloud computing algorithms for discrete constrained potential games," *Proceedings of the 5th IFAC Workshop on Distributed Estimation and Control in Networked Systems*, vol. 48, no. 22, pp. 180–185, September 2015.

[14] ——, "Privacy preserving distributed optimization using homomorphic encryption," *Automatica*, vol. 96, no. 10, pp. 314–325, October 2018.

[15] 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 *Proceedings of the 2016 IEEE 55th Conference on Decision and Control*, December 2016, pp. 5053–5058.

[16] M. Ruan, H. Gao, and Y. Wang, "Secure and privacy-preserving consensus," *IEEE Transactions on Automatic Control. To appear*, 2019.

[17] P. Paillier, "Public-key cryptosystems based on composite degree residuosity classes," in *Proceedings of Advances in Cryptology, EUROCRYPT 1999*, 1999, pp. 223–238.

[18] O. Goldreich, *Foundations of Cryptography: Volume 2-Basic Applications*. Cambridge University Press, 2004.

[19] J. Lian, H. Ren, Y. Sun, and D. J. Hammerstrom, "Performance evaluation for transactive energy systems using double-auction market," *IEEE Transactions on Power Systems*, 2019, to appear.

[20] C. Hazay and Y. Lindell, *Efficient Secure Two-Party Protocols–Techniques and Constructions*. Springer, 2010.

[21] *Weather Underground: weather record for Columbus*. [Online]. Available: https://www.wunderground.com/

[22] W. Marion and K. Urban, "User's manual for TMY2s: Typical meteorological years: Derived from the 1961–1990 national solar radiation data base," Golden, CO, USA: National Renewable Energy Lab, Tech. Rep., 1995.

[23] *PJM wholesale market energy price*. [Online]. Available: http://pjm.com/markets-and-operations/energy.aspx

[24] *AEP Ohio power company standard tariff*. [Online]. Available: https://aepohio.com/account/bills/rates/AEPOhioRatesTariffsOH.aspx

[25] *GridLAB-D Residential Module User's Guide*. [Online]. Available: http://www.eps.ee.kth.se/personal/luigiv/pst/