

A General Privacy-Preserving Auction Mechanism for Secondary Spectrum Markets

Qianyi Huang, Yang Gui, *Student Member, ACM*, Fan Wu, *Member, IEEE*, Guihai Chen, *Member, IEEE*, and Qian Zhang, *Fellow, IEEE*

Abstract—Auctions are among the best-known market-based tools to solve the problem of dynamic spectrum redistribution. In recent years, a good number of strategy-proof auction mechanisms have been proposed to improve spectrum utilization and to prevent market manipulation. However, the issue of privacy preservation in spectrum auctions remains open. On the one hand, truthful bidding reveals bidders' private valuations of the spectrum. On the other hand, coverage/interference areas of the bidders may be revealed to determine conflicts. In this paper, we present PISA, which is a Privacy preserving and Strategy-proof Auction mechanism for spectrum allocation. PISA provides protection for both bid privacy and coverage/interference area privacy leveraging a privacy-preserving integer comparison protocol, which is well applicable in other contexts. We not only theoretically prove the privacy-preserving properties of PISA, but also extensively evaluate its performance. Evaluation results show that PISA achieves good spectrum allocation efficiency with light computation and communication overheads.

Index Terms—Mechanism design, privacy preservation, radio spectrum management.

I. INTRODUCTION

INDUSTRY experts indicate that the fast-growing wireless technology is being stalled by the scarcity of radio spectrum [10]. This scarcity is often considered to be a result of

Manuscript received June 30, 2014; revised January 08, 2015 and April 13, 2015; accepted May 10, 2015; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor D. Leith. This work was supported in part by grants from the 973 project 2012CB316201, 2013CB329006; the China NSF under Grants 61173156, 61422208, 61472252, 61272443, and 61133006; the CCF-Intel Young Faculty Researcher Program and CCF-Tencent Open Fund; the Scientific Research Foundation for the Returned Overseas Chinese Scholars; the Jiangsu Future Network Research under Project No. BY2013095-1-10; the RGC under the contracts CERG 622613, 16212714, HKUST6/CRF/12R, M-HKUST609/13; and the Huawei-HKUST joint lab under a grant. The opinions, findings, conclusions, and recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agencies or the government. (*Corresponding author: Fan Wu.*)

Q. Huang is with the Department of Computer Science and Engineering, Hong Kong University of Science and Technology, Hong Kong, and also with the Shanghai Key Laboratory of Scalable Computing and Systems, Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai 200240, China (e-mail: qhuangaa@ust.hk).

Y. Gui, F. Wu, and G. Chen are with the Shanghai Key Laboratory of Scalable Computing and Systems, Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai 200240, China (e-mail: yanggui1989@gmail.com; fwu@cs.sjtu.edu.cn; gchen@cs.sjtu.edu.cn).

Q. Zhang is with the Department of Computer Science and Engineering, Hong Kong University of Science and Technology, Hong Kong (e-mail: qianzh@cs.ust.hk).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNET.2015.2434217

the static and rigid spectrum allocation by the government. The spectrum may be idle when the primary users are not engaged in data transmission, while at the same time many unlicensed users are starving for radio spectrum. Such a static allocation mechanism cannot fully utilize the limited spectrum. In order to improve spectrum utilization, secondary spectrum markets have emerged, where auctions are used to dynamically redistribute channels (e.g., [2], [8], [9], [12], [13], [33], [36], [37], and [46]). Different from the auctions held by the government, auctions in secondary spectrum markets occur dynamically and more frequently. The auctioneers may be primary users who tend to lease their channels in order to receive proper payoff during their idle time. The bidders may be secondary wireless service providers that need spectrum to serve their subscribers, or a mobile device that needs spectrum to transmit data.

In spectrum auctions, strategy-proofness (defined in Section II) is the topic of major research efforts, which stimulates bidders to bid their true valuations of the spectrum. It eliminates the overhead of gaming over each other, and the auctioneer can allocate the channels to who value it the most. However, different from the primary spectrum auctions where all bids are open and auction results are posted on FCC Web pages, there are privacy concerns in secondary spectrum auctions. Truthful bidding divulges the bidder's true valuations toward the spectrum, which are closely related to the profits of winning the spectrum. Bidders are not willing to share such information with other bidders or the auctioneer. Let us consider that the auctioneer is a cellular network provider, named A , and another cellular network provider B participates in the auction as a bidder to request for channels to transmit data. The bid may imply B 's economic situation, which is highly sensitive information. B is reluctant to disclose it to the auctioneer (i.e., A), who is a competitor in some sense. Furthermore, corrupt auctioneers may exploit such knowledge to their advantage. For instance, the auctioneers/sellers may set the reserve price accordingly in future auctions to increase their own revenue. Unfortunately, most existing works fail to protect bid privacy in auction design.

Moreover, spectrum allocation may disclose the bidders' coverage/interference areas. Spectrum is spatially reusable. Two bidders distanced by space can simultaneously use the same channel for transmission. In auctions, bidders may be required to reveal their coverage/interference areas to the auctioneer to determine conflict. However, coverage/interference areas may divulge the location information of the bidders, especially when the bidders are mobile devices. It may also disclose other sensitive information, such as their business models and subscriber




Fig. 1. System architecture.

distribution. Thus, bidders are reluctant to share their coverage/interference information with the auctioneer.

Therefore, privacy preservation and strategy-proofness are both important factors in designing spectrum auctions. However, there are several challenges. First, due to the spatial reusability of spectrum, well-separated bidders can share the same channel. Existing privacy-preserving auctions (e.g., [4] and [28]) are designed for traditional goods (e.g., paintings, jewelry), where each commodity can only be allocated to one bidder. When it comes to spectrum auctions, they may either fail or lead to significant degradation of spectrum utilization. Second, strategy-proofness and bid privacy are somewhat contradictory objectives. Strategy-proofness encourages bidders to reveal their true valuations of the spectrum, while bid privacy tends to prevent the auctioneer and other participants from learning the bidders' true valuations. Third, different from conventional auctions, spectrum allocation is constrained by geographic conditions. The allocation process should satisfy the geographic constraints while preserving bidders' coverage/interference area privacy.

In this paper, we consider the problem of privacy preservation in spectrum auctions and propose PISA, which is a Privacy preserving and Strategy-proof Auction mechanism for secondary spectrum markets. As shown in Fig. 1, we introduce a third party (e.g., [21], [22], and [32]), namely an agent, who acts as an intermediary between the bidders and the auctioneer. The agent should be nonprofit, and we require the agent to be a well-established organization. Therefore, some trustworthy nonprofit organizations are suitable to play the role of the agent, such as Spectrum Bridge [26]. Although the agent may be a well-established party, bidders are still reluctant to share private information with any party, the agent being no exception. Thus, in PISA, the agent and the auctioneer cooperate to perform the auction, but neither of them can infer any sensitive information about the bidders without collusion. The essence of PISA lies in our privacy-preserving bid comparison protocol, and we further extend the protocol to privately determine geographic conflict.

We summarize our contributions as follows.

- 1) To the best of our knowledge, PISA is the first strategy-proof spectrum auction mechanism that protects both bid privacy and coverage/interference area privacy without sacrificing social welfare.
- 2) We present a protocol to perform efficient comparison between integers without revealing their actual values. Our protocol can compare arbitrary large integers and is well applicable in other contexts.
- 3) We implement PISA and extensively evaluate its performance. The evaluation results show that PISA achieves

good channel utilization, with low computation and communication overhead.

The rest of the paper is organized as follows. In Section II, we review technical preliminaries. In Section III, we present the detailed design of basic PISA, which preserves bid privacy of the winners. In Section IV, we enhance PISA to provide stronger privacy protection (i.e., coverage/interference area privacy and k -anonymous bid privacy for all bidders). We present our evaluation results in Section V. In Section VI, we briefly review related works. Finally, we conclude and point out future research directions in Section VII.

II. PRELIMINARIES

In this section, we first briefly review some solution concepts and present our auction model. Then, we define a generic strategy-proof spectrum auction mechanism. Finally, we introduce a useful homomorphic cryptosystem.

A. Solution Concepts

We recall the solution concepts used in our study. Let s_i denote player i 's preference strategy and s_{-i} denote the strategy profile of all the players except for player i . $u_i(s_i, s_{-i})$ is the utility of player i when the strategy of player i is s_i and the strategies of all other players are s_{-i} .

Definition 1 (Incentive-Compatible [23]): A mechanism is incentive-compatible if for any strategy $s'_i \neq s_i$ and any other players' strategy profile s_{-i} , the utility u_i of the player i always satisfies the following condition:

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}).$$

Intuitively, in an incentive-compatible mechanism, players can maximize their utilities by reporting truthful preference information, regardless of other players' strategy profiles. Next, we introduce another related concept.

Definition 2 (Individual Rational [23]): For any strategy s_i and any other players' strategy profile s_{-i} , the utility u_i of the player i always satisfies the following condition:

$$u_i(s_i, s_{-i}) \geq 0.$$

A mechanism is individually rational if each player always gets a nonnegative utility, which means that each player can gain no less utility from faithful participation than nonparticipation. We now give the formal definition of *strategy-proof mechanism*.

Definition 3 (Strategy-Proof Mechanism [19], [30]): A mechanism is strategy-proof when it satisfies both incentive-compatibility and individual-rationality.

In a strategy-proof mechanism, misbehavior cannot result in any extra profit. Each player can maximize her utility by truthful participation.

In the field of privacy preservation, k -anonymity [29] is widely used to quantify the degree of privacy preservation (e.g., [39]). A scheme provides k -anonymous protection when a person cannot be distinguished from at least $k - 1$ other individuals.

Definition 4 (k -anonymity [29]): A privacy-preserving scheme satisfies k -anonymity if a participant cannot be

identified by the sensitive information with probability higher than $1/k$.

B. Auction Model

We model the procedure of spectrum allocation as a sealed-bid auction, involving an auctioneer, an agent, and a number of bidders. For clarity, we assume that there is a single channel to be shared among the bidders.¹

The auctioneer may be a primary user who tends to lease her idle channel to a group of conflict-free secondary users in order to receive proper payoff during her idle time. The auctioneer may also be a specialized third-party platform for spectrum management, such as Spectrum Bridge [26]. Bidders simultaneously submit their encrypted bidding tuples via the agent to the auctioneer. The auctioneer decides the winners and their charges.

We consider that there is a set $\mathbb{B} = \{1, 2, \dots, z\}$ of bidders. Let $\vec{v} = (v_1, v_2, \dots, v_z)$ denote the valuation profiles of the bidders, which is their private information. Accordingly, the bidders' bidding profile is denoted by $\vec{b} = (b_1, b_2, \dots, b_z)$. Let $\vec{x} = (x_1, x_2, \dots, x_z)$ and $\vec{y} = (y_1, y_2, \dots, y_z)$ denote the vector of bidders' latitudes and longitudes, respectively. Bidder i can share the channel with bidder j , if their coverage/interference areas do not overlap.

The auctioneer determines the charging profile $\vec{p} = (p_1, p_2, \dots, p_z)$ and the allocation profile $\vec{a} = (a_1, a_2, \dots, a_z)$, where $a_i = 1$ indicates that bidder i is allocated the channel, while $a_i = 0$ indicates not. The utility of bidder i can be defined as

$$u_i = (v_i - p_i)a_i.$$

The goal of all bidders is to maximize their own utilities. Here, we assume that the bidders do not collude with each other.

In contrast to the bidders, the overall objective of the auction is to guarantee strategy-proofness and to maximize channel allocation efficiency subject to the conflicting conditions.

C. Generic Spectrum Auction Scheme

As pointed out in [40], the spatial interference constraints make the problem of finding the optimal allocation in the above auction model NP-complete. A practical solution is to resort to monotonic allocation in order to improve computation efficiency and to apply critical charging to guarantee strategy-proofness. In this section, we present a representative auction scheme that achieves both strategy-proofness and computational efficiency.

We model the geographic conflicts of the bidders as a conflict graph \mathbb{G} . On the conflict graph, each bidder is represented by a vertex. Two bidders (vertices) are connected if their coverage/interference areas overlap. Two important concepts are *critical neighbor* and *critical value*.

Definition 5 (Critical Neighbor [40]): Given $\mathbb{G} \setminus \{i\}$, the critical neighbor $CN(i)$ of bidder i is a neighbor of i where if

¹For multichannel case, we can change the monotonic allocation algorithm and the critical charging algorithm according to [40] and apply our privacy-preserving bid comparison protocol.

i bids lower than $CN(i)$, i will not be allocated, and if i bids higher than $CN(i)$, i will be allocated.

Definition 6 (Critical Value [40]): The critical value of bidder i is defined as the bid of $CN(i)$; if $CN(i)$ does not exist, the critical value of i is 0.

Algorithm 1 describes the monotonic allocation procedure, where \mathbb{A} denotes the set of available bidders and $N(i)$ denotes the set of neighbors of bidder i in \mathbb{G} . Lines 2–6 iteratively allocate the channel to the highest bidder i in \mathbb{A} , and eliminate i and her neighbors from further consideration. The algorithm stops when there is no bidder left in \mathbb{A} .

Algorithm 1: Monotonic Allocation Algorithm

Input: Conflict graph \mathbb{G} , bidder set \mathbb{B} , and bidding profile \vec{b} .
Output: Allocation profile \vec{a} .
1: $\mathbb{A} \leftarrow \mathbb{B}; \vec{a} \leftarrow 0^z$.
2: **while** $\mathbb{A} \neq \emptyset$ **do**
3: $i \leftarrow \arg \max_{j \in \mathbb{A}} (b_j)$.
4: $a_i \leftarrow 1$.
5: $\mathbb{A} \leftarrow \mathbb{A} \setminus (N(i) \cup \{i\})$.
6: **end while**
Return \vec{a} .

Algorithm 2 shows the critical charging procedure. Lines 5–11 determine the critical value of bidder i . Each winner is charged with his/her critical value. In the case of a tie, we may break the tie either randomly or by the bidders' identifiers.

Algorithm 2: Critical Charging Algorithm

Input: Conflict graph \mathbb{G} , bidder set \mathbb{B} , bidding profile \vec{b} , allocation profile \vec{a} , and bidder i .
Output: Payment p_i .
1: **if** $a_i = 0$ **then**
2: **Return** 0.
3: **end if**
4: $\mathbb{A} \leftarrow \mathbb{B} \setminus \{i\}$.
5: **while** $\mathbb{A} \neq \emptyset$ **do**
6: $k \leftarrow \arg \max_{j \in \mathbb{A}} (b_j)$.
7: **if** $k \in N(i)$ **then** $\triangleright CN(i) = k$.
8: **Return** b_k .
9: **end if**
10: $\mathbb{A} \leftarrow \mathbb{A} \setminus (N(k) \cup \{k\})$.
11: **end while**
Return 0.

Here we have the first theorem. Please refer to [40] for the proof.

Theorem 1: The generic spectrum auction mechanism is strategy-proof.

We note that the generic spectrum auction mechanism may reveal the bidders' private information to the auctioneer. We present our approaches to protect bidders' privacy in Sections III and IV.

D. Boneh–Goh–Nissim (BGN) Cryptosystem

Homomorphic encryption is a form of encryption that enables specific types of computations to be carried out on ciphertexts and obtain a new ciphertext, which can be decrypted to match the result of computations applied directly to the original plaintexts.

In this work, we adopt Boneh–Goh–Nissim (BGN) cryptosystem [3]. It supports computations of unlimited number of additions with at most one multiplication. Thus, it can evaluate quadratic multivariate polynomials on ciphertexts.

Before introducing the BGN cryptosystem, we recall *bilinear map* and *bilinear group*, which are the bases of the BGN cryptosystem.

Definition 7 (Bilinear Map [3]): Let \mathbb{G}_1 and \mathbb{G}_2 be two cyclic groups of order n , for some large n . A map $e : \mathbb{G}_1 \times \mathbb{G}_1 \rightarrow \mathbb{G}_2$ is said to be bilinear if

$$e(P^a, Q^b) = e(P, Q)^{ab}$$

for all $P, Q \in \mathbb{G}_1$ and $a, b \in \mathbb{Z}$.

Definition 8 (Bilinear Group [3]): \mathbb{G}_1 is a bilinear group if there exists a group \mathbb{G}_2 and a bilinear map e , s.t.:

- 1) \mathbb{G}_1 and \mathbb{G}_2 are two multiplicative cyclic groups of finite order n .
- 2) g is a generator of \mathbb{G}_1 .
- 3) $e : \mathbb{G}_1 \times \mathbb{G}_1 \rightarrow \mathbb{G}_2$ is a bilinear map and $e(g, g)$ is a generator of \mathbb{G}_2 .

Given n , [3] presents an approach to constructing a bilinear group of order n . Due to limitations in space, we do not elaborate on it here. Next, we describe the three algorithms making up the BGN cryptosystem.

KeyGen(τ): Given a security parameter $\tau \in \mathbb{Z}^+$, generate two random τ -bit primes q_1, q_2 , and set $n = q_1 q_2 \in \mathbb{Z}$. Generate a bilinear group \mathbb{G}_1 of order n as described in [3]. Let $e : \mathbb{G}_1 \times \mathbb{G}_1 \rightarrow \mathbb{G}_2$ be the bilinear map.

Randomly pick two generators g, u from \mathbb{G}_1 and set $h = u^{q_2}$. Then, h is a random generator of a q_1 -order subgroup of \mathbb{G}_1 . The public key is $\mathcal{PK} = (n, \mathbb{G}_1, \mathbb{G}_2, e, g, h)$. The private key is $\mathcal{SK} = q_1$.

Encrypt(\mathcal{PK}, m): We assume that the message space consists of integers from $\{0, 1, \dots, T\}$ with $T < q_2$. To encrypt a message m with the public key \mathcal{PK} , pick a random integer r from \mathbb{Z}_n , and compute

$$C = g^m h^r \in \mathbb{G}_1.$$

Here, C is the ciphertext of m .

Decrypt(\mathcal{SK}, C): To decrypt a ciphertext C using the private key $\mathcal{SK} = q_1$, note that

$$C^{q_1} = (g^m h^r)^{q_1} = (g^{q_1})^m.$$

To recover m , it suffices to compute the discrete log of C^{q_1} in base g^{q_1} . Although it appears inefficient to do decryption, BGN cryptosystem is well suited to our scenario. In our application, we only need to decide whether a ciphertext is an encryption of 0 or not.

Next, we show the homomorphic properties of BGN cryptosystem. Given ciphertexts $C_1 = g^{m_1} h^{r_1} \in \mathbb{G}_1$ and $C_2 = g^{m_2} h^{r_2} \in \mathbb{G}_1$

$$C_1 \oplus C_2 = C_1 C_2 h^r = g^{m_1} h^{r_1} g^{m_2} h^{r_2} \cdot h^r = g^{m_1+m_2} h^{r_1+r_2+r}$$

is the ciphertext of $m_1 + m_2$ for a random $r \in \mathbb{Z}_n$.

Furthermore, BGN cryptosystem allows one multiplication using the bilinear map. Since h is of order q_1 , we rewrite $h = g^{\alpha q_2}$ for some unknown $\alpha \in \mathbb{Z}$. Set $\hat{g} = e(g, g)$ and $\hat{h} = e(g, h) = \hat{g}^{\alpha q_2}$. Hence, \hat{g} is of order n and \hat{h} is of order q_1 . Pick a random $r' \in \mathbb{Z}_n$, then

$$\begin{aligned} C_1 \otimes C_2 &= e(C_1, C_2) \hat{h}^{r'} = e(g^{m_1} h^{r_1}, g^{m_2} h^{r_2}) \hat{h}^{r'} \\ &= e(g^{m_1 + \alpha q_2 r_1}, g^{m_2 + \alpha q_2 r_2}) \hat{h}^{r'} \\ &= e(g, g)^{(m_1 + \alpha q_2 r_1)(m_2 + \alpha q_2 r_2)} \hat{h}^{r'} \\ &= e(g, g)^{m_1 m_2} e(g, g)^{\alpha q_2 (m_1 r_2 + m_2 r_1 + \alpha q_2 r_1 r_2)} \hat{h}^{r'} \\ &= e(g, g)^{m_1 m_2} e(g, h)^{m_1 r_2 + m_2 r_1 + \alpha q_2 r_1 r_2} \hat{h}^{r'} \\ &= \hat{g}^{m_1 m_2} \hat{h}^{\bar{r}} \in \mathbb{G}_2 \end{aligned}$$

is the ciphertext of $m_1 \times m_2$, where $\bar{r} = m_1 r_2 + m_2 r_1 + \alpha q_2 r_1 r_2 + r'$. We note that the system is still additively homomorphic in \mathbb{G}_2 .

III. BASIC PISA

In this section, we first introduce our privacy-preserving bid comparison protocol. Then, we introduce our spectrum auction mechanism, namely PISA, which preserves the winners' bid privacy and achieves strategy-proofness. Here, we assume that the auctioneer has full knowledge of bidders' coverage/interference areas to construct the conflict graph. In Section IV, we enhance our design to provide stronger privacy protection, that is, coverage/interference area privacy and k -anonymous bid privacy for both winners and losers.

A. Privacy-Preserving Bid Comparison

To determine the auction winners, it suffices to let the auctioneer know whether b_i is higher than b_j , for any i and j from \mathbb{B} , without revealing the exact values of b_i and b_j . This problem is a variant of secure comparison. A generic solution is based on Yao's garbled circuits [38], which have a predefined number of inputs. However, in spectrum auctions, it is difficult to determine the number of bidders before the bidding phase. Hence, Yao's approach is not practical here. Therefore, we design a more flexible approach. Our protocol is based on bitwise comparison, allowing us to compare arbitrary large integers. We describe it separately for clarity, and it is well applicable in other contexts.

We consider two l -bit binary bids $b_i = (b_i^l b_i^{l-1} \dots b_i^1)_2$ and $b_j = (b_j^l b_j^{l-1} \dots b_j^1)_2$, where b_i^l and b_j^l denote the least significant bits, while b_i^1 and b_j^1 denote the most significant bits. For each integer $k \in [1, l]$, we define

$$\omega_{ij}^k = (b_i^k - b_j^k)^2 = (b_i^k)^2 + (b_j^k)^2 - 2b_i^k b_j^k \quad (1)$$

$$\lambda_{ij}^k = \zeta_{ij}^k [(b_i^k)^2 - (b_j^k)^2 + 1 + \sum_{r=k+1}^l \omega_{ij}^r] \quad (2)$$

where $\zeta_{ij}^k \in_{\mathbb{R}} \mathbb{Z}^+$ is a random positive number. In this section, l, k , and t are all integers. Then, we get the following lemma.

Lemma 1: For any $i, j \in \mathbb{B}$, we have $b_i < b_j$, if and only if there exists exactly one $k \in [1, l]$, where $\lambda_{ij}^k = 0$.

Proof: For any $t \in [1, l]$, we have $\zeta_{ij}^t > 0$ and

$$\begin{aligned} \omega_{ij}^t &= (b_i^t - b_j^t)^2 \geq 0 \\ (b_i^t)^2 - (b_j^t)^2 + 1 &\geq 0 - 1 + 1 = 0 \end{aligned}$$

hence $\lambda_{ij}^t \geq 0$. Next, we prove the necessary and sufficient conditions.

- Given $b_i < b_j$, there exists $k \in [1, l]$, such that $b_i^k = 0 < b_j^k = 1$ and $b_i^t = b_j^t$ for any $t \in [k+1, l]$. Consequently

$$\omega_{ij}^k = (b_i^k - b_j^k)^2 = 1$$

and

$$\omega_{ij}^t = (b_i^t - b_j^t)^2 = 0, \forall t \in [k+1, l].$$

Hence

$$\lambda_{ij}^k = \zeta_{ij}^k [0^2 - 1^2 + 1 + 0] = 0.$$

We further distinguish two cases.

— For $t \in [1, k-1]$, since $\omega_{ij}^t = 1$, we have

$$\sum_{r=t+1}^l \omega_{ij}^r > 0.$$

Hence, $\lambda_{ij}^t > 0$.

— For $t \in [k+1, l]$

$$(b_i^t)^2 - (b_j^t)^2 + 1 = 1 > 0.$$

Hence, $\lambda_{ij}^t > 0$.

Therefore, there exists exactly one $k \in [1, l]$, where $\lambda_{ij}^k = 0$.

- Given $\lambda_{ij}^k = 0$, for $\zeta_{ij}^k > 0$, we can infer that

$$(b_i^k)^2 - (b_j^k)^2 + 1 = 0 \wedge \sum_{r=k+1}^l \omega_{ij}^r = 0.$$

Hence, $b_i^k = 0 < b_j^k = 1$ and $b_i^t = b_j^t$ for any $t \in [k+1, l]$.

Therefore, $b_i < b_j$.

This completes our proof. \square

We note that both (1) and (2) are quadratic polynomials in b_i^k and b_j^k . Consequently, we can evaluate them using the BGN cryptosystem. With Lemma 1, we can compare two bids without knowing their exact values.

B. Design Details

In basic PISA, bidders submit their encrypted bidding tuples to the agent who preprocesses them before transferring them to the auctioneer. The auctioneer can decrypt the encrypted tuples and find only the necessary information to run the auction, without inferring any nonessential information. In this section, we present the design details of basic PISA, which comprises four phases shown as follows.

Phase 1: Initialization: Before the auction, the auctioneer sets up the parameters for BGN cryptosystem and runs KeyGen(τ) (as shown in Section II-D). Then, the public key $\mathcal{PK} = (n, \mathbb{G}_1, \mathbb{G}_2, e, g, h)$ is announced, while the private key $\mathcal{SK} = q_1$ is not revealed. The auctioneer also sets the possible bidding range of integers $R = [b_{\min}, b_{\max}]$, where b_{\min} and b_{\max} are two l -bit binaries.

Phase 2: Bidding: In the bidding phase, each bidder i decides her bid $b_i \in R$ according to her valuation v_i . The bidder i encrypts every bit b_i^k from her bid b_i with the auctioneer's public key \mathcal{PK}

$$E(b_i^k) = \text{Encrypt}(\mathcal{PK}, b_i^k), k \in [1, l]$$

where $\text{Encrypt}()$ is the encryption function defined in Section II-D. For ease of expression, we denote the series of encrypted bits of b_i as

$$E(\tilde{b}_i) = (E(b_i^k))_{k \in [1, l]}.$$

Then, the bidder i sends $[i, E(\tilde{b}_i)]$ as her bidding tuple to the agent.

Phase 3: Preprocessing: After receiving all the encrypted bidding tuples from the bidders, the agent preprocesses the ciphertexts.

For each bidder i , the agent appends $E(\tilde{i})$ to the least significant end of $E(\tilde{b}_i)$. Here, similar with $E(\tilde{b}_i)$, $E(\tilde{i})$ is the series of encrypted bits of bidder i 's binary ID number. Now the bidding tuple turns out to be

$$[i, E(\tilde{b}_i) || E(\tilde{i})]$$

where $||$ is the concatenation operator. The suffix does not affect the comparison result of the two bids, except the case of tie. With the suffix, the tie can be broken according to the bidders' ID number.

Then, for any pair of bidders i and j , the agent computes

$$\begin{aligned} E(\omega_{ij}^k) &= E((b_i^k)^2 + (b_j^k)^2 - 2b_i^k b_j^k) \\ &= e(E(b_i^k), E(b_j^k)) \times e(E(b_j^k), E(b_j^k)) \\ &\quad \times e(E(b_i^k), E(b_j^k))^{-2} \\ E(\lambda_{ij}^k) &= E((b_i^k)^2 - (b_j^k)^2 + 1) + \sum_{r=k+1}^{l + \lceil \log_2 z \rceil} \omega_{ij}^r \zeta_{ij}^k \\ &= e(E(b_i^k), E(b_i^k)) \zeta_{ij}^k \times e(E(b_j^k), E(b_j^k))^{-\zeta_{ij}^k} \\ &\quad \times e(g, g)^{\zeta_{ij}^k} \times \prod_{r=k+1}^{l + \lceil \log_2 z \rceil} E(\omega_{ij}^r) \zeta_{ij}^k \end{aligned} \quad (3)$$

for each $k \in [1, l + \lceil \log_2 z \rceil]$ and $\zeta_{ij}^k \in_{\mathbb{R}} \mathbb{Z}^+$. Here, $\lceil \log_2 z \rceil$ is the length of bidders' binary ID number, since there are z bidders in total.

Finally, the bidder sends the following tuples to the auctioneer:

$$[i, j, E(\tilde{\lambda}_{ij})], \forall i, j \in \mathbb{B} \wedge i \neq j$$

where $E(\tilde{\lambda}_{ij})$ is the list of $E(\lambda_{ij}^k)$ for $k \in [1, l + \lceil \log_2 z \rceil]$. We note that the elements in $E(\tilde{\lambda}_{ij})$ can be randomly permuted.

Phase 4: Opening:

a) Conflict Graph Construction: The auctioneer can construct the conflict graph $\mathbb{G} = (\mathbb{V}, \mathbb{E})$ according to bidders' geographic distribution, where each bidder is represented by a vertex. Two vertices are connected if their coverage/interference areas overlap. Here, $N(f) = \{h \in \mathbb{V} | (f, h) \in \mathbb{E}\}$.

b) Monotonic Allocation: For each edge $(f, h) \in \mathbb{E}$, the auctioneer can decrypt $E(\tilde{\lambda}_{fh})$, and check whether it contains a λ_{fh}^k that is equal to 0 for $k \in [1, l + \lceil \log_2 z \rceil]$. If so, $b_f < b_h$; Otherwise, $b_f > b_h$.

We can update Algorithm 1 to protect the bidders' bidding values. Algorithm 3 shows the privacy-preserving winner allocation procedure. We define matrix $\mathbb{S} = [\tilde{\lambda}_{fh}]_{f, h \in \mathbb{V}}$. In lines 2–6, we iteratively pick bidder f , who does not have 0 in $\tilde{\lambda}_{fh}$, for any $h \in \mathbb{A} \setminus \{f\}$ (i.e., bidder f has the highest bid in \mathbb{A}); and eliminate f and her neighbors from \mathbb{A} . When the set \mathbb{A} becomes empty, the algorithm outputs the allocation profile \vec{a} .

Algorithm 3: Privacy-Preserving Allocation Algorithm

Input: Conflict graph (\mathbb{V}, \mathbb{E}) and matrix $\mathbb{S} = [\tilde{\lambda}_{fh}]_{f, h \in \mathbb{V}}$.

Output: Allocation profile \vec{a} .

```

1:  $\mathbb{A} \leftarrow \mathbb{V}$ ;  $\vec{a} \leftarrow 0^z$ .
2: while  $\mathbb{A} \neq \emptyset$  do
3:   Pick  $f \in \mathbb{A}$ , s.t.,  $\nexists h \in \mathbb{A} \setminus \{f\}$  and  $k \in [1, l + \lceil \log_2 z \rceil]$ ,
   such that  $\lambda_{fh}^k = 0$ .
4:    $a_f \leftarrow 1$ .
5:    $\mathbb{A} \leftarrow \mathbb{A} \setminus (N(f) \cup \{f\})$ .
6: end while
Return  $\vec{a}$ .

```

c) Critical Charging: Since the auctioneer is not given the encrypted bids from the agent, the charges to the winners cannot be computed directly according to Algorithm 2. However, the auctioneer can determine the critical neighbor of each winner.

Algorithm 4: Privacy-Preserving Critical Neighbor Determination Algorithm— $CN(f)$

Input: Conflict graph (\mathbb{V}, \mathbb{E}) , matrix $\mathbb{S} = [\tilde{\lambda}_{fh}]_{f, h \in \mathbb{V}}$, and bidder f .

Output: Critical neighbor $CN(f)$.

```

1:  $\mathbb{A} \leftarrow \mathbb{V} \setminus \{f\}$ .
2: while  $\mathbb{A} \neq \emptyset$  do
3:   Pick  $f' \in \mathbb{A}$ , s.t.,  $\nexists h \in \mathbb{A} \setminus \{f'\}$  and  $k \in [1, l + \lceil \log_2 z \rceil]$ ,
   such that  $\lambda_{f'h}^k = 0$ .
4:   if  $f' \in N(f)$  then
5:     Return  $f'$ .
6:   end if
7:    $\mathbb{A} \leftarrow \mathbb{A} \setminus (N(f') \cup \{f'\})$ .
8: end while
Return  $NULL$ .

```

Algorithm 4 shows our privacy-preserving critical neighbor determination procedure. In lines 2–8, we determine bidder f 's neighbor f' , who would be allocated the channel if f is absent

from the auction. Then, the algorithm outputs f' as the critical neighbor of f . If no such f' exists, the algorithm returns $NULL$. We note that Algorithm 4 differs from Algorithm 2 because it outputs the bidder f 's critical neighbor $CN(f)$ instead of the critical value.

d) Outcome Announcement: We denote the vector of critical neighbors by $\vec{CN} = (CN(f))_{f \in \mathbb{V}: a_f = 1}$. The auctioneer needs to resort to the agent for the encrypted bids of the critical neighbors. Next, the agent replies with a vector of encrypted bids of the critical neighbors

$$\vec{C} = (E(\tilde{b}_i))_{i \in \vec{CN}}.$$

Finally, the auctioneer can decrypt the encrypted critical bids in \vec{C} , and announce the winners together with their charges.

C. Analysis

We consider the computational complexity for the bidders, the agent and the auctioneer, respectively: Bidders have to carry out BGN encryption for each bit in their bids. Thus, each bidder has to carry out encryption l times, where l is the number of bits in the bid; the agent has to carry out preprocessing for each pair of bidders, so the computational complexity of the agent is $O(z^2l)$. Here, z is the number of bidders; both the allocation algorithm and charging algorithm run at $O(z^2l)$, thus the computational complexity of the auctioneer is $O(z^2l)$.

The strategy-proofness of PISA is inherited from the generic auction mechanism. We omit the proof here and directly draw the following conclusion.

Theorem 2: PISA satisfies strategy-proofness.

Besides privacy preservation and strategy-proofness, PISA also achieves the following nice properties.

- 1) Compared to the generic spectrum auction scheme, PISA protects bidders' privacy without sacrificing spectrum allocation efficiency.
- 2) In PISA, bidders are allowed to choose their bids from a contiguous integer range. Compared to existing mechanisms (e.g., [15]), in which bids are limited to a small set of predefined values, PISA provides bidders with more bidding flexibility.
- 3) PISA preserves the communication pattern of an auction protocol. In PISA, bidders are not required to communicate with each other. After submitting the bidding tuples to the agent, they are free from burdensome computation and communication.

Basic PISA is based on our privacy-preserving bid comparison protocol, hence the auctioneer can compare two bids without knowing their values. As an intermediary, the agent cannot decrypt the encrypted bidding tuples to learn the bids. Therefore, we protect the bid privacy of winners against both the auctioneer and the agent. However, the bids of critical neighbors are revealed as the charges of the winners. In the next section, we enhance our design to thwart such privacy breaches and provide protection for coverage/interference area privacy.

IV. EXTENDED PISA

In this section, we intend to provide k -anonymous bid privacy for both winners and losers, together with protection

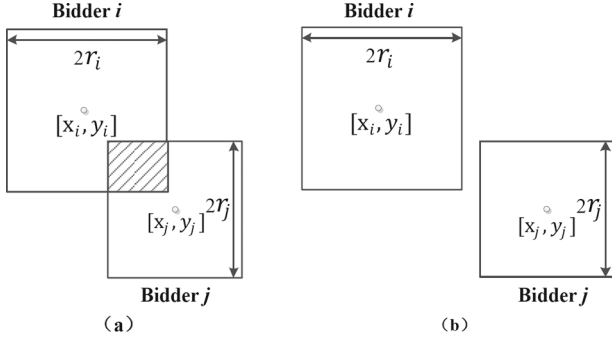


Fig. 2. Square coverage/interference area examples. In case (a), bidders i and j have conflict. In case (b), the two bidders do not have conflict.

for coverage/interference area privacy. To preserve coverage/interference area privacy, we generalize each bidder's coverage/interference area to a square with side length $2r_i$. As shown in Fig. 2, bidders i and j are conflicting bidders in case (a), while they can share the same channel in case (b). Compared to the commonly used circular conflict areas, this assumption may overestimate interference. To evaluate the impact of this assumption on channel allocation, in Section V, we compare our square conflict model to a conflict graph obtained from a real measurement.

We define

$$\begin{aligned} x_i^\triangleright &= x_i + r_i & x_i^\triangleleft &= x_i - r_i \\ y_i^\triangleup &= y_i + r_i & y_i^\triangleright &= y_i - r_i \end{aligned}$$

where x_i^\triangleright , x_i^\triangleleft , y_i^\triangleup , and y_i^\triangleright are t -bit binaries. Bidders i and j are out of the coverage/interference range of each other if the following condition holds:

$$(x_i^\triangleright < x_j^\triangleleft \vee x_i^\triangleleft > x_j^\triangleright) \vee (y_i^\triangleup < y_j^\triangleright \vee y_i^\triangleright > y_j^\triangleup) = TRUE. \quad (4)$$

Given Lemma 1, we can evaluate $x_i^\triangleright < x_j^\triangleleft$, $x_i^\triangleleft > x_j^\triangleright$, $y_i^\triangleup < y_j^\triangleright$, and $y_i^\triangleright > y_j^\triangleup$ privately. Due to limitations of space, we only focus on the differences in phases of bidding, preprocessing, and opening.

A. Design Details

Besides the processing of coverage/interference areas, the key difference lies in the preprocessing phase, where the agent performs a secret permutation to anonymize the bidders.

Phase 1: Initialization: It is similar to the basic PISA in Section III-B.

Phase 2: Bidding: In addition to $E(\tilde{b}_i)$, each bidder i also calculates $E(\tilde{x}_i^\triangleright)$, $E(\tilde{x}_i^\triangleleft)$, $E(\tilde{y}_i^\triangleup)$, and $E(\tilde{y}_i^\triangleright)$. Then, the bidder sends

$$[i, E(\tilde{b}_i), E(\tilde{x}_i^\triangleright), E(\tilde{x}_i^\triangleleft), E(\tilde{y}_i^\triangleup), E(\tilde{y}_i^\triangleright)]$$

to the agent as the bidding tuple.

Phase 3: Preprocessing: For any pair of bidders i and j , the agent computes $E(\tilde{\lambda}_{ij})$ as specified in Section III-B. In addition, the agent calculates

$$E(\alpha_{ij}^k) = E((x_i^{\triangleright k})^2 - (x_j^{\triangleleft k})^2 + 1 + \sum_{r=k+1}^t (x_i^{\triangleright r} - x_j^{\triangleleft r})^2)^{\xi_{ij}^k}$$

for each $k \in [1, t]$ and $\xi_{ij}^k \in_R \mathbb{Z}^+$. Same as before, we let $E(\tilde{\alpha}_{ij}) = (E(\alpha_{ij}^k))_{k \in [1, t]}$. The agent also carries out the calculations on $(x_i^\triangleleft, x_j^\triangleright)$, $(y_i^\triangleup, y_j^\triangleright)$ and $(y_i^\triangleright, y_j^\triangleup)$, and results in $E(\tilde{\beta}_{ij})$, $E(\tilde{\gamma}_{ij})$, and $E(\tilde{\delta}_{ij})$, respectively.

After finishing the above calculations, the agent carries out a secret permutation π on the results to make them anonymous to the auctioneer. Then, the agent sends the following anonymous tuples to the auctioneer:

$$[\pi(i), \pi(j), E(\tilde{\lambda}_{ij}), E(\tilde{\alpha}_{ij}), E(\tilde{\beta}_{ij}), E(\tilde{\gamma}_{ij}), E(\tilde{\delta}_{ij})]$$

for each $i, j \in \mathbb{B}$ and $i \neq j$.

Phase 4: Opening: The auctioneer decrypts $E(\tilde{\alpha}_{fh})$, $E(\tilde{\beta}_{fh})$, $E(\tilde{\gamma}_{fh})$, and $E(\tilde{\delta}_{fh})$, to get $\tilde{\alpha}_{fh}$, $\tilde{\beta}_{fh}$, $\tilde{\gamma}_{fh}$, and $\tilde{\delta}_{fh}$, respectively. We note that bidder f and h cannot share the channel, if the following condition holds:

$$\Omega_{fh} = (0 \notin \tilde{\alpha}_{fh} \wedge 0 \in \tilde{\beta}_{fh} \wedge 0 \notin \tilde{\gamma}_{fh} \wedge 0 \in \tilde{\delta}_{fh}) = TRUE. \quad (5)$$

The auctioneer can construct the conflict graph $\mathbb{G} = (\mathbb{V}, \mathbb{E})$, where

$$\begin{aligned} \mathbb{V} &= \{\pi(i) | i \in \mathbb{B}\} \\ \mathbb{E} &= \{(f, h) | f, h \in \mathbb{V} \wedge \Omega_{fh} = TRUE\}. \end{aligned}$$

Then, the auctioneer carries out Algorithms 3 and 4 on permuted bidders. Algorithm 3 outputs allocation profile \vec{a}' on permuted bidders, while Algorithm 4 returns the permuted critical neighbors. Since the auctioneer is unaware of the agent's perturbation $\pi: \mathbb{B} \rightarrow \mathbb{V}$, the original identifiers of the bidders remain anonymous. Let $\vec{W} = (f)_{f \in \mathbb{V}: a'_f = 1}$ denote the vector of permuted winners, and $\vec{CN} = (CN(f))_{f \in \mathbb{V}: a'_f = 1}$ denote the corresponding vector of permuted critical neighbors. To prevent the agent from finding one-to-one correspondence between the winners and their critical neighbors, the auctioneer permutes \vec{CN} to get \vec{CN}' . The auctioneer needs to resort to the agent for the winners' original identifiers and the encrypted bids of the critical neighbors, by sending \vec{W} and \vec{CN}' to the agent. Next, the agent replies with a vector of winner identifiers

$$\vec{W}' = (i)_{\pi(i) \in \vec{W}}$$

and a vector of encrypted bids of the critical neighbors

$$\vec{C}' = (E(\tilde{b}_i))_{\pi(i) \in \vec{CN}'}$$

Finally, the auctioneer can decrypt the encrypted critical bids in \vec{C}' , map the critical bids to the winners by reversing the permutation done on \vec{CN}' , and announce the winners together with their charges.

B. Illustrative Example

The following example may help to illustrate how extended PISA works. Suppose there are five bidders 1–5, located at (6, 10), (10, 14), (10, 6), (14, 10), and (10, 16), respectively. They are supposed to bid higher than 0, but lower than 16. Each of them has a valuation of the channel in binary form: $b_1 =$

$0010_2, b_2 = 0100_2, b_3 = 1010_2, b_4 = 0101_2, b_5 = 0110_2$. For simplicity, we assume $r_1 = r_2 = r_3 = r_4 = r_5 = 5$.

In the auction, each bidder submits an encrypted bidding tuple, which contains his/her encrypted bid and coverage/interference boundaries. The bidding tuple follows the format of $[i, E(\tilde{b}_i), E(\tilde{x}_i^{\triangleright}), E(\tilde{x}_i^{\triangleleft}), E(\tilde{y}_i^{\triangleleft}), E(\tilde{y}_i^{\triangleright})]$. For instance, bidder 1 submits

$$[1, E(\widetilde{0010}), E(\widetilde{1011}), E(\widetilde{0001}), E(\widetilde{1111}), E(\widetilde{0101})]$$

to the agent. Here, 1 is the identity

$$\begin{aligned} b_1 &= 0010_2 \\ x_1^{\triangleright} &= 6 + r_1 = 1011_2, x_1^{\triangleleft} = 6 - r_1 = 0001_2 \\ y_1^{\triangleleft} &= 10 + r_1 = 1111_2, y_1^{\triangleright} = 10 - r_1 = 0101_2. \end{aligned}$$

Similarly, bidder 3 submits

$$[3, E(\widetilde{1010}), E(\widetilde{1111}), E(\widetilde{0101}), E(\widetilde{1011}), E(\widetilde{0001})].$$

After collecting all the bids, the agent carries out preprocessing as specified in Phase 3. For simplicity, we omit the suffix. Then, the agent anonymizes all the tuples and sends the permuted ones to the auctioneer. For example, the agent sends the following information about bidders 1 and 3 to the auctioneer:

$$\begin{aligned} &[4'(3), 5'(1), E(\mathbf{2} \times \mathbf{2}, \mathbf{4} \times \mathbf{2}, \mathbf{1} \times \mathbf{2}, \mathbf{3} \times \mathbf{2}) \\ &E(\mathbf{4} \times \mathbf{2}, \mathbf{5} \times \mathbf{3}, \mathbf{3} \times \mathbf{4}, \mathbf{1} \times \mathbf{4}), E(\mathbf{3} \times \mathbf{0}, \mathbf{5} \times \mathbf{3}, \mathbf{3} \times \mathbf{2}, \mathbf{2} \times \mathbf{4}) \\ &E(\mathbf{7} \times \mathbf{2}, \mathbf{3} \times \mathbf{1}, \mathbf{2} \times \mathbf{4}, \mathbf{2} \times \mathbf{4}), E(\mathbf{3} \times \mathbf{0}, \mathbf{3} \times \mathbf{1}, \mathbf{2} \times \mathbf{2}, \mathbf{2} \times \mathbf{4})]. \end{aligned}$$

Here, the numbers inside parentheses in the first two terms are the bidders' true identifiers, which are hidden from the auctioneer. Furthermore, $\pi(1) = 5'$ and $\pi(3) = 4'$ are the anonymous identification after the secret permutation π . The numbers in bold are the random numbers generated by the agent [e.g., 2, 4, 1, 3 in the third term correspond to the random number ζ_{ij} in (3)].

The auctioneer can decrypt the ciphertexts and do the comparison. In this example

$$\begin{aligned} \Omega_{4'5'} &= [0 \notin (\mathbf{4} \times \mathbf{2}, \mathbf{5} \times \mathbf{3}, \mathbf{3} \times \mathbf{4}, \mathbf{1} \times \mathbf{4})] \\ &\wedge [0 \in (\mathbf{3} \times \mathbf{0}, \mathbf{5} \times \mathbf{3}, \mathbf{3} \times \mathbf{2}, \mathbf{2} \times \mathbf{4})] \\ &\wedge [0 \notin (\mathbf{7} \times \mathbf{2}, \mathbf{3} \times \mathbf{1}, \mathbf{2} \times \mathbf{4}, \mathbf{2} \times \mathbf{4})] \\ &\wedge [0 \in (\mathbf{3} \times \mathbf{0}, \mathbf{3} \times \mathbf{1}, \mathbf{2} \times \mathbf{2}, \mathbf{2} \times \mathbf{4})] \\ &= \text{TRUE}. \end{aligned}$$

Thus, bidders $4'$ and $5'$ cannot share the channel. The auctioneer constructs the conflict graph based on the comparison results, as shown in Fig. 3. The term by each edge $(f, h) \in \mathbb{E}$ denotes $[f, h, E(\tilde{\lambda}_{fh})]$ (elements in $E(\tilde{\lambda}_{fh})$ are multiplied with random positive numbers and are randomly perturbed).

From Fig. 3, the auctioneer can learn that

$$\begin{aligned} b_{1'} &> b_{2'}, \{b_{1'}, b_{3'}\} > b_{2'} > b_{5'} \\ b_{4'} &> b_{3'} > b_{2'}, b_{4'} > \{b_{3'}, b_{5'}\}, \{b_{2'}, b_{4'}\} > b_{5'}. \end{aligned}$$

Then, the auctioneer can determine winners and their critical neighbors using Algorithms 3 and 4, respectively. The winners turn out to be bidders $1'$ and $4'$, and the corresponding critical




Fig. 3. Conflict graph constructed by the auctioneer, without revealing actual coverage/interference areas of the bidders.

neighbors are bidders $2'$ and $3'$, respectively. The auctioneer permutes $\overrightarrow{CN} = \{2', 3'\}$ to get $\overrightarrow{CN}' = \{3', 2'\}$.

Next, the auctioneer consults the agent with $\overrightarrow{W} = \{1', 4'\}$ and $\overrightarrow{CN}' = \{3', 2'\}$, and the agent replies with answer 5, 3, $E(\tilde{b}_4)$, and $E(\tilde{b}_2)$. Finally, the auctioneer announces bidders 3 and 5 as winners, with charges 5 and 4, respectively.

C. Analysis

It is evident that extended PISA inherits the nice properties from basic PISA, including strategy-proofness and the three properties listed in Section III-C. To avoid redundancy, we do not elaborate on them again here. In this section, we demonstrate some other nice properties of our auction mechanism.

Theorem 3: Extended PISA guarantees k -anonymity for bid privacy, where $k = |\overrightarrow{CN}|$.

Proof: We distinguish the following two cases.

Case 1: Bidder i is a critical neighbor, i.e., $\pi(i) \in \overrightarrow{CN}$. We consider from the perspectives of both the agent and the auctioneer.

On the one hand, the agent cannot decrypt bidder i 's bidding tuple without the secret key SK , hence the agent cannot find out what the bid is. When the auctioneer consults the agent with \overrightarrow{CN}' in the opening phase, the agent learns that bidder i is one of the critical neighbors. However, the order of bidders in \overrightarrow{CN}' has been permuted by the auctioneer. The agent does not know which bidder is which winner's critical neighbor. When the auctioneer announces the charges for winners, bidder i 's bid is hidden among $|\overrightarrow{CN}'|$ charges. Thus, the agent cannot identify bidder i 's bid with probability higher than $1/|\overrightarrow{CN}'|$.

On the other hand, although the auctioneer can decrypt the bidding tuples, they cannot be linked with the bidders because all the bidding tuples are anonymized by the agent. The true identifier of a critical neighbor is hidden among $z - |\overrightarrow{W}|$ losers. Thus, the auctioneer cannot identify bidder i with probability higher than $1/(z - |\overrightarrow{W}|)$.

For the set of critical neighbors is a subset of losers ($z - |\overrightarrow{W}| \geq |\overrightarrow{CN}'|$), both the agent and the auctioneer cannot identify bidder i 's bid with probability higher than $1/|\overrightarrow{CN}'|$.

Case 2: Bidder i is not a critical neighbor, i.e., $\pi(i) \notin \overrightarrow{CN}$.

For bidder $\pi(i) \notin \overrightarrow{CN}$, the agent cannot decrypt the bidding tuple, and the bid is never revealed to the auctioneer. We note that with the suffix, all bids have different encrypted values. There is no possibility that the auctioneer can infer a bid as it happens to be equal to one of the critical values.

This completes our proof. \square

As for coverage/interference areas, the agent cannot decrypt bidders' encrypted bidding tuples and can only perform homomorphic operations on the ciphertexts. Although the auctioneer can decrypt the preprocessed tuples received from the agent, the bidders' coverage/interference areas remain unknown. The auctioneer constructs the conflict graph based on Lemma 1, without knowledge of the bidders' exact coverage/interference boundaries. Hence, bidders' coverage/interference areas are revealed to neither the auctioneer nor the agent. Therefore, we have the following theorem.

Theorem 4: Extended PISA prevents the agent and the auctioneer from learning the bidders' coverage/interference areas.

We also analyze the computational complexity: In extended PISA, besides bids, we have to process the four coverage/interference area boundaries. Each bidder has to carry out BGN encryption for $(l + 4t)$ times. Here, t is the number of bits in the four boundaries; similarly, the agent has to carry out preprocessing for each pair of bidders, so the computational complexity of the agent is $O(z^2(l + 4t))$; for each pair of bidders, the auctioneer has to judge whether their conflict squares overlap. Thus, the auctioneer has to carry out $(4z^2t)$ times of decryption to construct the conflict graph. The allocation algorithm and charging algorithm still run at $O(z^2l)$, thus the computational complexity of the auctioneer is $O(z^2(l + 4t))$.

V. EVALUATION

We have implemented PISA and evaluated its efficiency and overhead through simulations. In this section, we show our evaluation results.

A. Methodology

In our evaluations, we implement the BGN cryptosystem with security parameter $\tau = 80$, using the Stanford pairing-based cryptography library (PBC), which is a C library built on the GMP library to perform the mathematic operations underlying pairing-based cryptosystems.

In each set of evaluations, we vary a factor among bidder number, the size of the terrain, and the number of digits in a bid, while fixing the other two. The number of bidders varies from 20 to 200, and the bidders are randomly distributed in the terrain. The size of the terrain ranges from 256×256 to 2048×2048 m², such that the coordinates can be represented by 8–11 bits. The coverage/interference range is randomly selected from 50 to 150 m, and hence the mean value is 100 m, which is the transmission range of 802.11b. The bid of each bidder ranges from 2 to 1023, which can be represented by at most 10 bits. The default values for bidder number, the size of terrain, and the number of digits in a bid is 200, 2048×2048 m², 10, respectively.

We measure the following metrics in our evaluation:

- Channel utilization: the average number of bidders allocated to the channel;
- Computational overhead: the processing time required by each party to run the auction;
- Communication overhead: the size of data that must be sent to convey information.

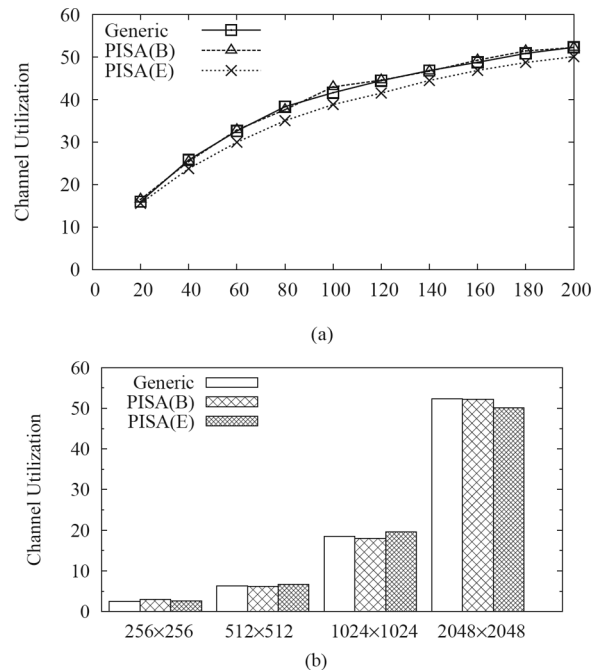


Fig. 4. Channel utilizations of the generic spectrum auction scheme, basic PISA (Section III), and extended PISA (Section IV). (a) Number of bidders. (b) Size of terrain area (meter \times meter).

We run a series of evaluations on a PC with Intel Core i5 3.1 GHz processor and 4 GB memory under Ubuntu 10. All the results on performance are averaged over 100 runs.

B. Allocation Efficiency

In this section, we compare PISA to VERITAS [40], i.e., the generic spectrum auction without privacy preservation.

Fig. 4(a) shows the channel utilizations achieved by the generic spectrum auction scheme, basic PISA, and extended PISA as a function of the number of bidders, when the terrain area is 2048×2048 m². We can see that the channel utilizations achieved by all the mechanisms are nondecreasing concave functions of the number of bidders.

Fig. 4(b) shows the case in which we vary the size of the terrain and fix the number of bidders at 200. Again, we can see that the generic auction scheme, basic PISA, and extended PISA all have increasing channel utilization. Bidders are randomly distributed over the terrain; a larger terrain results in less conflicts and hence higher channel utilization.

Fig. 4 validates our claim that, compared to the generic spectrum auction scheme, PISA protects bidders' privacy without sacrificing channel utilization.

To show the impact of square conflict area assumption, we compare it to a conflict graph obtained from real measurements. We utilize the data collected by Zhou *et al.* [41]. This dataset contains 78 APs of the Google WiFi network, covering a 7-km² residential area in Mountain View, CA, USA.

We obtain two conflict graphs: one from the real measurement [41], and the other one constructed by the square conflict model. Here, we set the interference range to be 150 m. The two conflict graphs are shown as follows. Fig. 5(a) is from [41], and Fig. 5(b) is constructed by the square conflict model. In

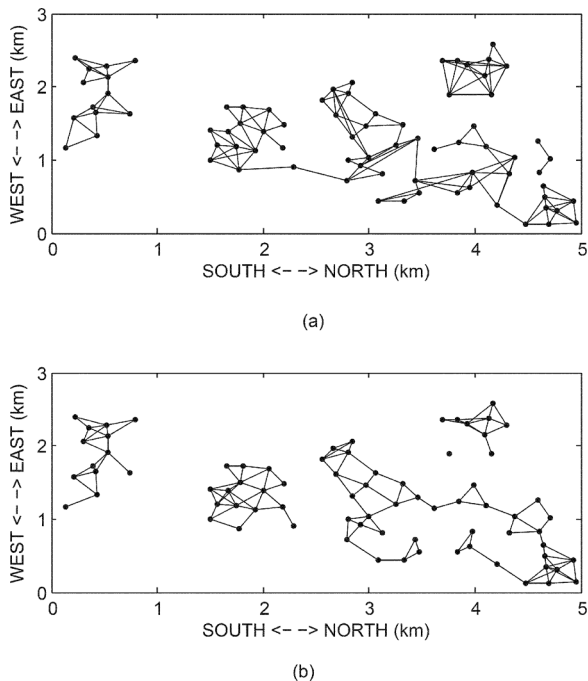


Fig. 5. Conflict graphs by (a) real measurement and (b) square conflict model.

Fig. 5(a), there are 151 edges, whereas in Fig. 5(b), there are 171 edges.

We use these two different conflict graphs as input and run the spectrum allocation algorithm. Bids are randomly distributed over $[50, 100]$. We repeat the experiments for 1000 runs. The average channel utilization for Fig. 5(a) is 26.2, whereas the number for Fig. 5(b) is 26.1. The two numbers are quite close. This is reasonable as there are only minor differences between the two conflict graphs. Thus, we can conclude that the assumption of the square conflict areas has a minor impact on channel utilization.

C. Overhead

PISA integrates cryptographic tools to protect bidders' privacy. An efficient privacy-preserving mechanism should have a low overhead. We evaluate the computation and communication overheads by varying the number of bidders, the size of the terrain, and the number of bits in the bid.

Fig. 6 shows the computational overhead of the agent and the auctioneer. We do not plot the bidders' computational overhead because each bidder just encrypts several bits of information and the computational overhead is only about 25 ms. We can see that the computational overhead is mainly from the agent because the agent is responsible for a large number of encryption operations. Furthermore, we can find that the agent spends far more time in extended PISA than in basic PISA. This is because in basic PISA, the agent spends most of the time processing bids. However, in extended PISA, the agent also needs to preprocess the four coverage/interference boundaries between bidders. Similarly, the auctioneer in extended PISA has to decrypt more ciphertexts to construct the conflict graph, hence the higher computational overhead than in basic PISA.

Specifically, Fig. 6(a) shows the run time against the number of bidders with 10-bit bids and a $2048 \times 2048\text{-m}^2$ terrain area.

We find that the computational overhead of the agent grows as a quadratic polynomial of bidders' number. This is reasonable because, in Phase 3: Preprocessing, the agent has to carry out preprocessing for each pair of bidders. As shown in Fig. 6(b), the computational overhead of the agent in basic PISA changes slightly as the size of terrain area grows. However, in extended PISA, the agent has to process the ciphertexts of coordinates for the auctioneer to build the conflict graph, thus her computational overhead increases with the size of the terrain. Fig. 6(c) shows that the run time of the agent increases almost linearly with the number of bits in a bid. This is reasonable, as for each pair of bidders, the agent has to compute (3) k times, where k is the number of bits in the bids. Generally, in our evaluations, the time required by the agent in basic PISA is less than 50 s, while in extended PISA, the agent requires about a few minutes' processing time (less than 5 min). To speed up computation, we can use parallel computing to save computation time. Since bidders are not involved in burdensome computation after submitting bids, they are not required to stay connected with the agent nor the auctioneer. They can simply wait for the auctioneer to broadcast the results, hence we believe this time gap is acceptable.

Fig. 7 plots the communication overhead induced by basic PISA and extended PISA. The communication overhead of each bidder is about 96 B in basic PISA and 546 B in extended PISA. It is trivial compared to the total communication overhead, hence we do not show them on the figures. As shown in Fig. 7, the communication overhead of extended PISA is much higher than basic PISA. This is because, in addition to transmitting the preprocessed results of bids, the agent in extended PISA has to transmit the preprocessed results of four coverage/interference boundaries to the auctioneer.

Similar to Fig. 6(a), the communication overhead in Fig. 7(a) grows as a quadratic polynomial of bidders' numbers. As shown in Fig. 7(b), the communication overhead of extended PISA grows almost linearly with the size of the terrain. With the increases in terrains, we need more bits to represent bidders' coverage/interference areas. Thus, communication overheads increase linearly with the number of bits needed to represent the terrain. Similar to the computational overhead, the communication overhead increases almost linearly with the number of bits in a bid, which is shown in Fig. 7(c). Generally, the communication overhead of basic PISA is less than 4 MB, while the communication overhead of extended PISA is less than 25 MB.

From Figs. 6 and 7, we can conclude that PISA protects bidders' privacy with tolerable computation and communication overheads. Since bidders are not engaged in burdensome computation and communication, both the computation and communication overheads for bidders are negligible, which is an appealing property in auction design.

VI. RELATED WORK

We briefly review related works in this section.

A. Privacy-Preserving Mechanism Design

Some works have been devoted to privacy-preserving mechanism design. Wang *et al.* [31] incentivized SUs to contribute




Fig. 6. Computational overheads induced by basic PISA (Section III) and extended PISA (Section IV). (a) Number of bidders. (b) Size of terrain area (meter \times meter). (c) Size of bid (bit).




Fig. 7. Communication overheads induced by basic PISA (Section III) and extended PISA (Section IV). (a) Number of bidders. (b) Size of terrain area (meter \times meter). (c) Size of bid (bit).

their sensing data for collaborative sensing by providing differential privacy protection in the presence of malicious service providers and SUs. Naor *et al.* [22] proposed Yao's garbled circuits for use in auctions. However, the number of bidders in spectrum auctions cannot be known before the bidding phase, hence Yao's garbled circuits are not applicable here. Sui and Boutilier [27] studied efficiency and privacy tradeoffs in mechanism design. Their results show that sacrifices in efficiency can provide gains in privacy. Similarly, Feigenbaum *et al.* [11] proposed a general framework to analyze the tradeoff between communication cost and privacy. In [7], the authors present a protocol based on homomorphic encryption for secure comparison of integers, which is well applicable for auctions. There are a great number of existing works on privacy-preserving auctions (e.g., [1], [5], [17], [25], and [28]), which are designated for traditional goods (e.g., paintings, jewelry), where each commodity can only be allocated to one bidder. When it comes to spectrum auctions, they may either fail or lead to significant degradation of spectrum utilization. Assume that we directly apply one of the existing privacy-preserving auction schemes to spectrum auctions, each channel will be allocated to only one bidder, which cannot fully exploit the spatial reusability of spectrum, resulting in extremely low channel utilization.

B. Dynamic Spectrum Auction

Auctions are widely used to handle spectrum allocation, and researchers have proposed various spectrum auction mechanisms (e.g., [35], [36], [40], and [43]). TAHES [12] and TRUST [42] are both truthful double spectrum auctions. Dong *et al.* [9] and Zhu *et al.* [46] applied combinatorial auctions to allocate spectrum. Deek *et al.* [8] and Xu *et al.* [37] investigated various forms of cheating in online auctions. Al-Ayyoub and Gupta [2] and Jia *et al.* [16] aimed at maximizing the revenue of primary users. Most of the existing literature mainly focuses on the economic aspects of the auction.

Pan *et al.* [24] proposed a secure spectrum auction leveraging paillier cryptosystem. Their design requires multiple auctioneers, which is normally considered to be impractical. Liu *et al.* [18] studied location privacy in spectrum auctions, however their auctions are not strategy-proof. In a closely related study, Huang *et al.* [15] proposed a novel spectrum auction mechanism to preserve bid privacy. However, SPRING [15] is based on bid-independent bidder grouping, and thus may result in terribly poor spectrum utilization in extreme cases. Furthermore, in [15], bidders are only allowed to choose bids from a small set of predefined values. PISA differs significantly from [15] as PISA is based on monotonic allocation and critical charging instead of bidder grouping. Moreover, PISA allows bidders to choose bids from a continuous integer range, which is more flexible. Furthermore, PISA protects both bid privacy and coverage/interference area privacy. Recently, Zhu *et al.* [44], [45] extended the exponential mechanism in [20] and proposed the first differentially privacy-preserving spectrum auction with approximate revenue maximization, under the assumption that the auctioneer is trustworthy. However, as mentioned before, bidders are reluctant to share their bidding information with anyone else, including the auctioneer. Thus, this assumption is not always true. In other related works [6], [14], [34], the authors adopted the similar system architecture. They mainly focused on protecting bid privacy, but did not consider bidders' coverage/interference area privacy.

VII. CONCLUSION AND FUTURE WORK

In this paper, we have presented PISA, which is the first privacy-preserving and strategy-proof spectrum auction mechanism that can protect both bid privacy and coverage/interference area privacy, without sacrificing social welfare. PISA is based on our novel and efficient privacy-preserving integer comparison protocol, which can compare arbitrary large integers and is well applicable in other contexts. Analytical results have demonstrated PISA's privacy-preserving properties

and evaluation results have shown that PISA achieves good spectrum allocation efficiency, with light computation and communication overheads.

As for future work, it will be interesting to study potential attacks against our model. Yet another possible direction is to provide privacy protection for both buyers and sellers in double spectrum auctions.

REFERENCES

- [1] M. Abe and K. Suzuki, "M + 1-st price auction using homomorphic encryption," in *Proc. PKC*, Feb. 2002, pp. 115–124.
- [2] M. Al-Ayyoub and H. Gupta, "Truthful spectrum auctions with approximate revenue," in *Proc. IEEE INFOCOM*, Apr. 2011, pp. 2813–2821.
- [3] D. Boneh, E.-J. Goh, and K. Nissim, "Evaluating 2-DNF formulas on ciphertexts," in *Proc. TCC*, Feb. 2005, pp. 325–341.
- [4] F. Brandt, "How to obtain full privacy in auctions," *Int. J. Inf. Security*, vol. 5, no. 4, pp. 201–216, 2006.
- [5] F. Brandt and T. Sandholm, "On the existence of unconditionally privacy-preserving auction protocols," *Trans. Inf. Syst. Security*, vol. 11, no. 2, p. 6, 2008.
- [6] Z. Chen *et al.*, "Ps-TRUST: provably secure solution for truthful double spectrum auctions," in *Proc. IEEE INFOCOM*, Apr. 2014, pp. 1249–1257.
- [7] I. Damgård, M. Geisler, and M. Krøigaard, "Efficient and secure comparison for on-line auctions," in *Proc. ACISP*, Jul. 2007, pp. 416–430.
- [8] L. Deek, X. Zhou, K. Almeroth, and H. Zheng, "To preempt or not: tackling bid and time-based cheating in online spectrum auctions," in *Proc. IEEE INFOCOM*, Apr. 2011, pp. 2219–2227.
- [9] M. Dong, G. Sun, X. Wang, and Q. Zhang, "Combinatorial auction with time-frequency flexibility in cognitive radio networks," in *Proc. IEEE INFOCOM*, Mar. 2012, pp. 2282–2290.
- [10] FCC, "Mobile broadband: The benefits of additional spectrum," FCC Staff Technical Paper, Oct. 2010.
- [11] J. Feigenbaum, A. D. Jaggard, and M. Schapira, "Approximate privacy: foundations and quantification (extended abstract)," in *Proc. EC*, Jun. 2010, pp. 167–178.
- [12] X. Feng, Y. Chen, J. Zhang, Q. Zhang, and B. Li, "TAHES: truthful double auction for heterogeneous spectrums," in *Proc. IEEE INFOCOM*, Mar. 2012, pp. 3076–3080.
- [13] L. Gao, Y. Xu, and X. Wang, "Map: multiauctioneer progressive auction for dynamic spectrum access," *IEEE Trans. Mobile Comput.*, vol. 10, no. 8, pp. 1144–1161, Aug. 2011.
- [14] H. Huang, X.-Y. Li, Y. e Sun, H. Xu, and L. Huang, "PPS: privacy-preserving strategyproof social-efficient spectrum auction mechanisms," *IEEE Trans. Parallel Distrib. Syst.*, vol. 26, no. 5, pp. 1393–1404, May 2014.
- [15] Q. Huang, Y. Tao, and F. Wu, "SPRING: a strategy-proof and privacy preserving spectrum auction mechanism," in *Proc. IEEE INFOCOM*, Apr. 2013, pp. 827–835.
- [16] J. Jia, Q. Zhang, Q. Zhang, and M. Liu, "Revenue generation for truthful spectrum auction in dynamic spectrum access," in *Proc. MobiHoc*, May 2009, pp. 3–12.
- [17] H. Lipmaa, N. Asokan, and V. Niemi, "Secure vickrey auctions without threshold TRUST," *Financial Cryptography*, vol. 2357, pp. 87–101, 2003.
- [18] S. Liu, H. Zhu, R. Du, C. Chen, and X. Guan, "Location privacy preserving dynamic spectrum auction in cognitive radio network," in *Proc. IEEE ICDCS*, Jul. 2013, pp. 256–265.
- [19] A. Mas-Colell, M. D. Whinston, and J. R. Green, *Microeconomic Theory*. Oxford, U.K.: Oxford Univ. Press, 1995.
- [20] F. McSherry and K. Talwar, "Mechanism design via differential privacy," in *Proc. IEEE FOCS*, Oct. 2007, pp. 94–103.
- [21] J. Meyerowitz and R. R. Choudhury, "Hiding stars with fireworks: Location privacy through camouflage," in *Proc. MobiCom*, Sep. 2009, pp. 345–356.
- [22] M. Naor, B. Pinkas, and R. Sumner, "Privacy preserving auctions and mechanism design," in *Proc. EC*, Oct. 1999, pp. 129–139.
- [23] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*. Cambridge, U.K.: Cambridge Univ. Press, 2007.
- [24] M. Pan, J. Sun, and Y. Fang, "Purging the back-room dealing: secure spectrum auction leveraging paillier cryptosystem," *IEEE J. Sel. Areas Commun.*, vol. 29, no. 4, pp. 866–876, Apr. 2011.
- [25] K. Sako, "An auction protocol which hides bids of losers," in *Proc. PKC*, Jan. 2000, pp. 422–432.
- [26] Spectrum Bridge, "Spectrum Bridge," [Online]. Available: <http://www.spectrumbridge.com>
- [27] X. Sui and C. Boutilier, "Efficiency and privacy tradeoffs in mechanism design," in *Proc. AAAI*, Aug. 2011, pp. 738–744.
- [28] K. Suzuki and M. Yokoo, "Secure combinatorial auctions by dynamic programming with polynomial secret sharing," *Financial Cryptography*, vol. 2357, pp. 44–56, 2003.
- [29] L. Sweeney, "K-anonymity: a model for protecting privacy," *Int. J. Uncertainty, Fuzziness Knowl.-Based Syst.*, vol. 10, no. 5, pp. 557–570, 2002.
- [30] H. Varian, "Economic mechanism design for computerized agents," in *Proc. USENIX Workshop Electron. Commerce*, 1995, p. 2.
- [31] W. Wang and Q. Zhang, "Privacy-preserving collaborative spectrum sensing with multiple service providers," *IEEE Trans. Wireless Commun.*, vol. 14, no. 2, pp. 1011–1019, Feb. 2015.
- [32] W. Wang and Q. Zhang, "Towards long-term privacy preservation: a context-aware perspective," *IEEE Wireless Commun.*, 2015, to be published.
- [33] X. Wang *et al.*, "Spectrum sharing in cognitive radio networks—an auction-based approach," *IEEE Trans. Syst., Man, Cybern. B, Cybern.*, vol. 40, no. 3, pp. 587–596, Jun. 2010.
- [34] F. Wu, Q. Huang, Y. Tao, and G. Chen, "Towards privacy preservation in strategy-proof spectrum auction mechanisms for noncooperative wireless networks," *IEEE/ACM Trans. Netw.*, 2014, DOI: 10.1109/TNET.2014.2322104, to be published.
- [35] F. Wu and N. Vaidya, "SMALL: a strategy-proof mechanism for radio spectrum allocation," in *Proc. IEEE INFOCOM*, Apr. 2011, pp. 81–85.
- [36] P. Xu, X.-Y. Li, S. Tang, and J. Zhao, "Efficient and strategyproof spectrum allocations in multichannel wireless networks," *IEEE Trans. Comput.*, vol. 60, no. 4, pp. 580–593, Apr. 2011.
- [37] P. Xu, X. Xu, S. Tang, and X.-Y. Li, "Truthful online spectrum allocation and auction in multi-channel wireless networks," in *Proc. IEEE INFOCOM*, Apr. 2011, pp. 26–30.
- [38] A. C. Yao, "Protocols for secure computations (extended abstract)," in *Proc. FOCS*, Nov. 1982, pp. 160–164.
- [39] H. Zang and J. Bolot, "Anonymization of location data does not work: A large-scale measurement study," in *Proc. MobiCom*, Sep. 2011, pp. 145–156.
- [40] X. Zhou, S. Gandhi, S. Suri, and H. Zheng, "Ebay in the sky: Strategy-proof wireless spectrum auctions," in *Proc. MobiCom*, Sep. 2008, pp. 2–13.
- [41] X. Zhou *et al.*, "Practical conflict graphs for dynamic spectrum distribution," in *Proc. ACM SIGMETRICS*, Jun. 2013, pp. 5–16.
- [42] X. Zhou and H. Zheng, "TRUST: a general framework for truthful double spectrum auctions," in *Proc. IEEE INFOCOM*, Apr. 2009, pp. 999–1007.
- [43] X. Zhou and H. Zheng, "Breaking bidder collusion in large-scale spectrum auctions," in *Proc. MobiHoc*, Sep. 2010, pp. 121–130.
- [44] R. Zhu, Z. Li, F. Wu, K. G. Shin, and G. Chen, "Differentially private spectrum auction with approximate revenue maximization," in *Proc. MobiHoc*, Aug. 2014, pp. 185–194.
- [45] R. Zhu and K. G. Shin, "Differentially private and strategy-proof spectrum auction with approximate revenue maximization," in *Proc. IEEE INFOCOM*, 2015, to be published.
- [46] Y. Zhu, B. Li, and Z. Li, "Core-selecting combinatorial auction design for secondary spectrum markets," in *Proc. IEEE INFOCOM*, Apr. 2013, pp. 1986–1994.



Qianyi Huang received the B.S. degree in computer science from Shanghai Jiao Tong University, Shanghai, China, in 2013, and is currently pursuing the Ph.D. degree in computer science and engineering at Hong Kong University of Science and Technology, Hong Kong.

Her research interests lie in privacy preservation and resource management in wireless networking.



Yang Gui is a graduate student with the Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai, China.

His research interests lie in mobile social network and resource management in wireless networking.

Mr. Gui is a student member of the Association for Computing Machinery (ACM) and CCF.



Fan Wu (M'10) received the B.S. degree in computer science from Nanjing University, Nanjing, China, in 2004, and the Ph.D. degree in computer science and engineering from the State University of New York at Buffalo, Buffalo, NY, USA, in 2009.

He is an Associate Professor with the Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai, China. He has visited the University of Illinois at Urbana-Champaign (UIUC), Urbana, IL, USA, as a Post-Doctoral Research Associate. He has published more than 80 peer-reviewed

papers in leading technical journals and conference proceedings. His research interests include wireless networking and mobile computing, algorithmic game theory and its applications, and privacy preservation.

Dr. Wu has served as the Chair of CCF YOCSEF Shanghai, on the Editorial Board of *Computer Communications*, and as the member of technical program committees of more than 40 academic conferences. He is a recipient of the China National Natural Science Fund for Outstanding Young Scientists, CCF-Intel Young Faculty Researcher Program Award, and CCF-Tencent "Rhinoceros bird" Open Fund, and is a Pujiang Scholar.



Guihai Chen (M'05) received the B.S. degree in computer science from Nanjing University, Nanjing, China, in 1984, the M. Engineering degree in computer engineering from Southeast University, Nanjing, China, in 1987, and the Ph.D. degree in computer science from the University of Hong Kong, Hong Kong, in 1997.

He visited Kyushu Institute of Technology, Kitakyushu, Japan, in 1998 as a Research Fellow, and the University of Queensland, Brisbane, Australia, in 2000 as a Visiting Professor. During 2001 to 2003, he was a Visiting Professor with Wayne State University, Detroit, MI, USA. He is a Distinguished Professor and Deputy Chair with the Department of Computer Science, Shanghai Jiao Tong University, Shanghai, China. He has published more than 200 papers in peer-reviewed journals and refereed conference proceedings in the areas of wireless sensor networks, high-performance computer architecture, peer-to-peer computing, and performance evaluation.

Prof. Chen is a member of the IEEE Computer Society. He has also served on technical program committees of numerous international conferences.



Qian Zhang (M'00–SM'04–F'12) received the B.S., M.S., and Ph.D. degrees in computer science from Wuhan University, Wuhan, China, in 1994, 1996, and 1999, respectively.

She joined Hong Kong University of Science and Technology, Hong Kong, in 2005, where she is now a Full Professor with the Department of Computer Science and Engineering. Before that, she was with Microsoft Research Asia, Beijing, China, in 1999, where she was the Research Manager of the Wireless and Networking Group. She has published about

300 refereed papers in international leading journals and key conferences in the areas of wireless/Internet multimedia networking, wireless communications and networking, wireless sensor networks, and overlay networking. Her current research is on cognitive and cooperative networks, dynamic spectrum access and management, as well as wireless sensor networks.

Dr. Zhang is a Fellow of IEEE for contribution to the mobility and spectrum management of wireless networks and mobile communications. She has received the MIT TR100 (*MIT Technology Review*) World's Top Young Innovator Award. She also received the Best Asia Pacific (AP) Young Researcher Award elected by the IEEE Communication Society in 2004.