

Privacy-Preserving Crowdsourced Spectrum Sensing

Xiaocong Jin
Arizona State University
Email: xiaocong.jin@asu.edu

Yanchao Zhang
Arizona State University
Email: yczhang@asu.edu

Abstract—Crowdsourced spectrum sensing has great potential in improving current spectrum database services. Without strong incentives and location privacy protection in place, however, mobile users will be reluctant to act as mobile crowdsourcing workers for spectrum sensing tasks. In this paper, we present PriCSS, the first framework for a crowdsourced spectrum sensing service provider to select spectrum-sensing participants in a differentially privacy-preserving manner. Thorough theoretical analysis and simulation studies show that PriCSS can simultaneously achieve differential location privacy, approximate social cost minimization, and truthfulness.

I. INTRODUCTION

Dynamic spectrum access (DSA) is an emerging paradigm for mitigating worldwide wireless spectrum shortage. A DSA system consists of licensed primary users and unlicensed secondary users. A secondary user can use a licensed channel currently not used by its primary user. With DSA in place, secondary users have more channels to use, and primary users can profit by sharing their under-utilized licensed spectrums.

Avoiding harmful interference with primary users is the first principal in DSA systems. FCC advocates a solution based on spectrum databases, each currently administrated by private entities such as Google and Microsoft. Each spectrum database administrator accepts registrations from primary users and leverages a well-known propagation model to predict the coverage boundary of each primary user. Each secondary user needs to inquire the spectrum database about the channel occupancy at a chosen location before transmitting there. Current spectrum databases have well-known drawbacks [1]. First, the signal propagation models in use are not accurate, leading to either severe under-utilization of the spectrum or interference with primary users. Second, current spectrum databases cannot provide the quality information of channels, which can significantly vary in space and time. Last, the locations of primary and secondary users cannot be validated, so a spectrum database administrator may return wrong spectrum occupancy information to secondary users.

Crowdsourced spectrum sensing (CSS) is very promising for mitigating the drawbacks of the current spectrum databases. In this approach, a spectrum database administrator recruits distributed mobile users to sense a given channel around a specified location and decides the channel occupancy by aggregating the sensing results. The feasibility of CSS is backed up by a few trends. First, the number of mobile devices are expected to hit 10 billion in 2016, which implies sufficient geographic coverage especially in populated metropolitan areas where DSA systems are expected to play significant roles. Second, future mobile devices are very likely to be capable of spectrum sensing given the expected pervasiveness of DSA-based wireless systems [2]. Last, mobile devices are

increasingly powerful in self-localization, communication, and computation, which has fostered the explosive popularity of mobile crowdsourcing applications [3]. With CSS in place, the spectrum database administrator does not need to deploy a dedicated large-scale sensor network for spectrum sensing.

A typical CSS system works as follows. The spectrum database administrator publishes spectrum-sensing tasks either periodically or randomly. Each spectrum-sensing task involves one or multiple channels, a pre-determined set of geographic locations, and the sensing time. The sensing results from the designated locations can be aggregated to jointly determine the channel occupancy at the specified time. Each mobile user in the CSS system can independently decide his capability of performing the sensing tasks. For example, if Tom will go to a restaurant for lunch today, he can conveniently perform the sensing task near the restaurant around the lunch time. Given the participating requests, the spectrum database administrator can select a set of users for each sensing task.

There are many challenges for pushing the promising CSS system above into practice. For example, strong incentives must be provided to stimulate self-interested mobile users for spectrum sensing. Incentive mechanism design for CSS systems is a non-trivial task. On the one hand, different users may want different rewards for the same sensing task. For instance, a user far away from the allocated location may require more to compensate for his longer driving time and higher fuel consumption; a user may also lie about his travel distance to a specific sensing location to gain more. On the other hand, the spectrum database administrator wants to minimize the overall participants' cost (i.e., social cost) for any sensing task as long as the sensing quality is sufficient. Another significant challenge lies in the location privacy of mobile users. Since spectrum-sensing tasks involve rich spatiotemporal information, the whereabouts of participating users can be easily exposed, thus discouraging mobile users wary of their location privacy.

This paper presents PriCSS, a novel framework for a spectrum database administrator to select spectrum-sensing participants in a differentially privacy-preserving manner. Our specific contributions are as follows. First, we formulate participant selection in CSS systems as a reverse auction problem where each participant's true cost for performing the sensing tasks is closely tied with the participant's current location. Second, we demonstrate a location-privacy attack under the previous formulation. Third, we present a new formulation based on the exponential mechanism to offer differential location privacy. Last, we thoroughly evaluate PriCSS through theoretical and simulation studies. Our results confirm that PriCSS can simultaneously achieve the following objectives.

- **Differential location privacy.** PriCSS can prevent any internal or external attacker with arbitrary knowledge from inferring the locations of mobile participants.
- **Approximate social cost minimization.** Social cost is the sum of the real cost of participants completing all the sensing tasks [4]. PriCSS aims to approximately minimize the social cost.
- **Truthfulness.** Each PriCSS participant has no incentive to lie about his sensing cost.

II. RELATED WORK

This section reviews the prior work most related to PriCSS.

There are a few elegant schemes on location privacy in CSS systems [5]–[8]. The majority of the schemes focus on preventing the spectrum database administrator from inferring the physical sensing locations based on submitted sensing reports. The authors in [8] introduce a framework for protecting location privacy of workers participating in spatial crowdsourcing tasks. In our context, the sensing locations are pre-determined and publicly known. PriCSS seeks to hide the current locations of sensing participants when competing to participate in the spectrum-sensing tasks, thus we try to address a very different problem.

Some other schemes aim to detect false sensing reports [9]–[14] or spectrum misuse [15]–[18]. PriCSS focuses on the pre-sensing phase and is orthogonal to these nice efforts.

Numerous efforts [4], [19]–[21] have been made on incentive mechanism design for crowdsourcing worker selection. Our work differs from this line of works by specifically addressing spectrum sensing and also location privacy.

Differential privacy [22]–[24] has been recently introduced into DSA research. The work in [25], [26] targets differentially private spectrum auctions. The work in [27] applies differential privacy to stream monitoring. In contrast, our work targets CSS systems and differential location privacy.

III. SYSTEM AND ADVERSARY MODELS

A. System Model

PriCSS is run by a spectrum database administrator whose functionalities, however, go far beyond those of the current spectrum database administrators. Specifically, similar to a spectrum database administrator, the PriCSS administrator accepts registrations from primary users and answers the spectrum-occupancy queries from secondary users. In addition, the PriCSS administrator can manage the spectrum of itself or other licensed users by issuing spatiotemporal spectrum permits which allow secondary users to use specific channels at specific locations during specific periods.

The PriCSS administrator relies on mobile crowdsourcing to obtain fine-grained information for its managed spectrum. Crowdsourcing spectrum sensing tasks eliminate the need for the PriCSS administrator to deploy and manage a large-scale sensor network dedicated to spectrum sensing. More specifically, to determine the realtime quality and occupancy of a specific channel in a certain area, the PriCSS administrator recruits mobile users there, referred to as PriCSS participants,

to perform spectrum sensing at a set of designated locations. The PriCSS administrator can then make a decision by fusing the sensing reports. This sensing method is known as cooperative spectrum sensing and has been widely studied. The sensing locations usually should be far apart from each other to ensure high spatial diversity and thus high sensing quality. For the purpose of this paper, we hereby assume that the PriCSS administrator has pre-determined the sensing locations of each sensing task according to the existing methods such as [28].

Each PriCSS participant is a mobile user who owns an advanced mobile device capable of spectrum sensing. He registers with the PriCSS administrator under his real identity to receive rewards for performing spectrum sensing. Each PriCSS participant also has a unique pseudonym or identifier which is visible to other participants in the system. In contrast, the real identity of each participant is kept confidential to himself or the PriCSS administrator.

B. Adversary Model

We assume that the PriCSS administrator is fully trusted in preserving the real identity and bids of PriCSS participants. This common assumption can be relaxed by introducing multiple semi-trusted parties who do not collude. How this relaxation can be done is beyond the scope of this paper.

The adversary can be internal or external to PriCSS. An internal attacker corresponds to a PriCSS participant. We assume that internal attackers are honest-but-curious (HBC) in the sense that they faithfully fulfill promised sensing tasks but have interests in finding out the locations of other PriCSS participants. We also assume that PriCSS participants may lie about their spectrum sensing cost to claim more rewards, but they are rational in the sense that they only lie if they can benefit. Such HBC and rational assumptions are commonly adopted in the literature to model the attackers not performing denial-of-service attacks. In contrast, an external attacker does not participate in PriCSS but tries to infer the locations of PriCSS participants from public information.

We assume that the adversary has arbitrary background knowledge for attempting to breach the location privacy. For example, both internal and external attackers know the details of the system operations, and they may also collude. We intend to offer differential location privacy to each PriCSS participant under this strong adversary model.

As mentioned in Section II, there can be many other security and privacy issues in CSS systems. We resort to the rich literature for effective defenses, e.g., detecting fake sensing results [9]–[14] and spectrum misuse [15]–[18].

IV. PARTICIPANT SELECTION WITHOUT PRIVACY

We first formulate participant selection in PriCSS as a reverse-auction problem without considering location privacy. For this purpose, we assume that there are totally n PriCSS participants in a large geographic region such as the Los Angeles metropolitan area. Each participant has a unique integer index in $\mathcal{N} = \{1, \dots, n\}$, which corresponds to his system pseudonym in practice.

We assume that the PriCSS administrator issues K sensing tasks. Each task $k \in [1, K]$ contains one or more channels to sense, a time window in which the sensing should be done, and $\mu_k \geq 1$ sensing locations which are determined by the PriCSS administrator according to existing results such as [28]. Finally, we denote the j -th subtask of task k by $t_{k,j}$, all the μ_k subtasks of task k by $T_k = \{t_{k,j} | j \in [1, \mu_k]\}$, and all the $\sum_{k=1}^K \mu_k$ subtasks by $\mathcal{T} = \{t_{k,j} | k \in [1, K], j \in [1, \mu_k]\}$. The participant is allowed to include multiple sensing tasks at once in his sensing request according to his schedule and itinerary. Since all the subtasks for the same sensing task need to be performed in the same (and generally short) time window, we require that each participant can at most perform one subtask for each sensing task.

The cost for spectrum sensing is modeled as follows. The PriCSS administrator publishes a constant factor η to compensate each PriCSS participant for his resource (power, communication, and computation) consumption and human effort incurred for each sensing subtask. Another constant ρ is also published as the travel compensation per unit distance for gas consumption, driving time, etc. For simplicity, we use Euclidean distance to model the travel distance between two points. Assume that a participant chooses to perform m subtasks in a round trip of total Euclidean distance d . His true sensing cost is defined as $v = m\eta + \rho d$. For example, if a participant is currently at position l_1 and he wants to perform two subtasks a and b which are located at l_a and l_b , respectively. Then d equals $\text{Euclidean}(l_1, l_a) + \text{Euclidean}(l_a, l_b) + \text{Euclidean}(l_b, l_1)$. Therefore, his true sensing cost for the two subtasks is simply $2\eta + \rho d$. Each participant knows this cost model for computing his sensing cost, and the PriCSS administrator can modify the model based on user feedbacks.

The PriCSS administrator aims to select n_k unique participants for each spectrum sensing task $k \in [1, K]$. Since PriCSS participants compete to perform spectrum sensing tasks in return for rewards, it is reasonable to model participant selection in PriCSS under a reverse combinatorial auction framework [29]. In this framework, the PriCSS administrator serves as an auctioneer to auction the sensing tasks, and each participant $i \in [1, n]$ acts as a bidder for the sensing tasks.

We outline the auction procedure as follows. The PriCSS administrator broadcasts the subtask list \mathcal{T} and expects each interested participant i to reply with one bid $b_i = (L_i, c_i)$, where $L_i \subset \mathcal{T}$, and c_i is his claimed cost to perform the sensing subtasks L_i . We assume that c_i is limited in the range of $[c_{\min}, c_{\max}]$, where c_{\min} and c_{\max} are reasonable minimum and maximum possible sensing costs, respectively. Each participant follows two rules to place his bid. First, he can bid for no more than one subtask for each sensing task. Second, he can bid for multiple sensing tasks. The first rule is necessary to prevent strategic manipulation of the bids. For example, participant A and B both bid for the same two subtasks $t_{1,1}$ and $t_{1,2}$. If bidding truthfully, A will be allocated with $t_{1,1}$ and B will be allocated $t_{1,2}$. However, A might find out that if he is assigned with $t_{1,2}$, he can gain more rewards. Thus, A could purposely lie about the cost of $t_{1,1}$ to give away

the sensing opportunity of $t_{1,1}$ to B . Since B has already been assigned with one subtask for this specific sensing task, B is excluded for consideration of task assignment of $t_{1,2}$. In this way, A purposely lies about one sensing cost to win the other sensing subtask and gains more. The second rule is to allow participants to perform multiple spectrum sensing tasks during a round trip so that the total cost for performing the bundled sensing tasks can be reduced.

Given the bid set $\mathcal{B} = \{b_i | i \in [1, n]\}$, the administrator determines the outcome of the auction, denoted by $\vec{x}(\mathcal{B}) = \{x_1, x_2, \dots, x_n\}$, where x_i is an indicator for participant i :

$$x_i = \begin{cases} 1, & i \text{ wins the subtask bundle } L_i, \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

Correspondingly, the administrator selects a winner set \mathcal{W} such that all subtasks in \mathcal{T} can be fulfilled.

Each participant also holds a true valuation about the performing cost for the subtask set L_i , which is calculated with the cost model previously and denoted by v_i . The utility of participant i whose bid b_i is accepted is defined as " $u_i = p_i x_i - v_i$ ", where p_i is the payment the administrator makes to participant i . Note that the utility is normalized to 0 if the participant is not a winner. The participants know the allocation algorithm and the payment scheme in advance, and each participant wants to choose his strategy to maximize his own utility. So the claimed cost c_i might not necessarily equal v_i for each participant.

In our model, for each sensing task bundle, the participants could have different valuations due to different sensing and travel cost involved. Since asking for what bundle is up to the participant to decide, we aim to design a truthful mechanism so that participants have no interests in lying about the claimed cost. In addition, the time interval between consecutive rounds of auctions can be dynamically adjusted by the PriCSS administrator according to his service requirements.

Problem Formulation. We formulate participation selection in PriCSS as follows without considering location privacy.

$$\begin{aligned} & \text{minimize} && \sum_{i \in \mathcal{W}} c_i \\ & \text{subject to} && |(\bigcup_{i \in \mathcal{W}} L_i) \cap T_k| = \mu_k, \forall k \in [1, K], \\ & && |L_i \cap T_k| \leq 1, \forall k \in [1, K], \forall i \in \mathcal{W}. \\ & && |L_i| \leq \gamma, \forall i \in \mathcal{W}. \end{aligned} \quad (2)$$

The first condition in the equation above indicates that participants in the winner set can fulfill all the K sensing tasks. The second one requires that each participant bid at most one subtask for each sensing task. The third one is to limit the number of sensing tasks a participant can perform in a single round. γ is a constant and specified by the administrator.

The basic problem can be essentially treated as a minimum weighted set cover problem [30], which is knowingly NP-hard. So our basic problem is also NP-hard, which can be solved by an iterative approximation algorithm as follows. We define the average contributory cost of a participant as his original

claimed cost over the number of subtasks which he bids for and are not yet allocated to other participants. In each iteration, the PriCSS administrator selects a new participant who has the minimum contributory cost among the remaining participants. The algorithm terminates when all the constraints are satisfied. We say that one participant *outbids* another if the former is chosen earlier than the latter.

V. YOUR LOCATION IS NO SECRET

In this section, we exemplify some attacks to infer PriCSS participants' locations when they are selected under the reverse auction framework in Section IV. The location of a participant here refers to his *base location* (e.g., home or workplace) where he stays for a long time each day, and the base location serves as the reference point for the participant to derive his sensing cost for any interested spectrum sensing tasks. We assume that each participant starts from his base location and returns there after performing spectrum sensing tasks. This assumption is reasonable in practice. For example, a service guy of a heating and air conditioning company always starts from and returns to his company after handling a sequence of service appointments; a company employee always returns to his workplace after having lunch; and a person always starts from and returns to his home at the end of the day.

We also assume that the PriCSS administrator publicizes each spectrum-sensing auction result to ensure the public that its participant selection is unbiased. The publicized information only includes the system identifier of each participant winning one or multiple sensing subtasks. The real identity, claimed cost, and received payment of each winning participant are still kept confidential. Making the auction result public can also help the winners achieve greater self-esteem and public recognition, for which there are numerous examples in practice. For instance, an Amazon user can get his product reviews seen and voted by others, and those contributing highly voted reviews can get free products to test and keep.

The key insight for the location-inference attacks is that a participant's claimed sensing cost is tied to his round-trip Euclidean distance according to the aforementioned public cost model $v = m\eta + \rho d$, which corresponds to performing m subtasks in a round trip of total Euclidean distance d . Even if the claimed cost of each participant is hidden, the attackers can still infer the locations of some participants from the auction results and the changes in auction participation. We give some attack examples in what follows to highlight the need for preserving location privacy. We consider two rounds of auctions, which involve identical channels and sensing locations but different sensing times. This is practical because the PriCSS administrator may want to know the occupancy and quality of each channel in each service area according to a periodic, on-demand, or random schedule.

Case 1: Single Task.

We first consider a simple case in which each participant can bid for a single sensing task. Since each participant can perform no more than one sensing subtask for any sensing task, the bid of each participant is hence for a single subtask.

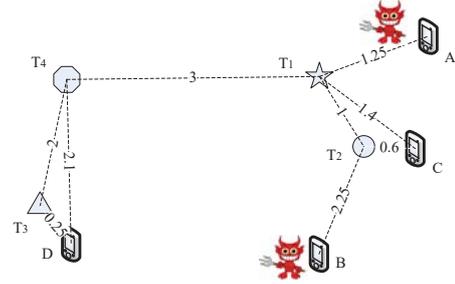


Fig. 1: A location-inference attack example.

For example, consider three participants $\{A, B, C\}$ bidding for the same subtask. According to the aforementioned cost model, their true sensing costs are $v_A = \eta + \rho d_A$, $v_B = \eta + \rho d_B$, and $v_C = \eta + \rho d_C$, respectively, where d_A , d_B , and d_C denote their respective Euclidean distance to the subtask location. Assume that the base locations of A , B , and C do not change. Nor do d_A , d_B , and d_C . In addition, we temporarily assume that the claimed cost of each participant equals his true sensing cost, which can be technically guaranteed later. So we have $c_A = v_A$, $c_B = v_B$, and $c_C = v_C$. Assuming that $d_A > d_B > d_C$, we have $c_A > c_B > c_C$. According to our formulation in Eq. (2), participant C will be selected as the winner in the first round. In the second round (say, next day), assuming that C no longer competes for this subtask for some reason such as work schedule change, so only A and B bid. Then B wins in the second round. The PriCSS administrator publishes the participant selection result in each round.

An external attacker can infer from the public information that $c_A > c_B > c_C$ and hence $d_A > d_B > d_C$, which are something a sensitive user does not want to disclose.

Internal attackers can infer much more information. For example, assume that B is an attacker. Since B knows his own distance d_B and $d_C < d_B$, he can infer that participant C must be inside the *suspicion region*, which is the circle centered at the subtask location with radius d_B . If C additionally participates in other sensing subtasks whose locations are also public, B can draw other suspicion regions for C and infer that C is in the intersection area of the suspicion regions with an overwhelming probability. B can also speed up his inference and improve the inference accuracy by colluding with other participants in the PriCSS system.

Case 2: Multiple Tasks.

We also give a more complicated example corresponding to the more general case in Eq. (2), in which each participant can bid for multiple subtasks with a single claimed cost. As shown in Fig. 1, our example involves four sensing tasks $T_1 \sim T_4$, each involving a single subtask. So we can use $T_1 \sim T_4$ to denote the four subtasks as well. The number associated with each dotted line in Fig. 1 represents the Euclidean distance between the two end locations. Let η be 0.5 and ρ be 1 for the aforementioned cost model $v = m\eta + \rho d$, where m denotes the number of chosen subtasks, and d denotes the round-trip Euclidean distance. The bids submitted by $A \sim D$ are as follows: $b_A = \{\{T_1\}, 3\}$, $b_B = \{\{T_2\}, 5\}$, $b_C = \{\{T_1, T_2\}, 4\}$, $b_D =$

$\{\{T_3, T_4\}, 5.35\}$. According to our formulation in Eq. (2), the winner set is $\mathcal{W} = \{C, D\}$. In the second round, assuming that C leaves the area or simply skips the auction, the winner set is $\mathcal{W}' = \{D, A, B\}$. Assume that the PriCSS administrator publishes a re-ordered winner set in each round to conceal each winner's selection order. For example, $\{D, C\}$ and $\{B, D, A\}$ are published as the two rounds' results.

There can be many attack strategies for the above scenario. Due to space limitations, we only discuss one case here, in which A and B collude to infer C 's location. The attack involves two steps. First, the attackers need to infer the sensing task bundle that C bids. Second, the attackers estimate the claimed cost of C . The first step can be achieved by studying the difference between the two winner sets, \mathcal{W} and \mathcal{W}' . From the attackers' point of view, D 's bid must have covered only T_3 and T_4 . Otherwise, the winner set would have been changed. It follows that C 's bid must have covered at least T_1 and T_2 . The remaining question is whether C 's bid also covers either or both T_3 and T_4 .

There are two possible cases now. In the first case, we assume that D outbids C in the first auction and thus gets T_3 and T_4 , so C can only contribute to tasks T_1 and T_2 . Since C outbids both A and B , his average contributory cost should be smaller than the smallest of A and B 's average contributory cost, which corresponds to $c_C/2 < c_A = 3$ or $c_C < 6$. From Fig. 1, the minimum round-trip cost for C to perform T_2 , T_1 and T_4 sequentially must be larger than 6 and is incurred when C first visits the T_2 location, then the T_1 location and the T_4 location, and finally C 's location. The additional cost is higher if T_3 is involved. So C 's bid covers T_1 and T_2 only. Plugging $m = 2$, $\eta = 0.5$, and $\rho = 1$ into the cost model $c_C = m\eta + \rho d_C$, the attackers have $c_C = 1 + d_C$ and thus $d_C < 5$. Since the distance between T_1 and T_2 is 1, the sum of the Euclidean distances from C to T_1 and T_2 is smaller than 4. So the attackers can infer that C must be inside the ellipse with T_1 and T_2 locations as two foci and the major-axis length equal to 4. C 's location can be further narrowed down if additional information is available.

Case 3.

In addition to the two exemplary attacks on location privacy, the participants very close to some subtask locations are likely to have lower claimed cost and higher chances to always win the sensing tasks at those locations, as the aforementioned approximate solution to our formulation in Eq. (2) is a deterministic process. Therefore, if a participant appears much more frequently than other participants in repeated auctions for the same sensing subtasks, the attackers can infer that the participant must be very close to one of the subtask locations. This kind of location privacy breach should also be prevented.

VI. PARTICIPANT SELECTION WITH DIFFERENTIAL LOCATION PRIVACY

Till now we have formulated participant selection in PriCSS as an NP-hard problem and described an approximate solution. We have also demonstrated a few attacks under the basic formulation and solution, which can severely endanger the

location privacy of PriCSS participants. In this section, we incorporate differential privacy into the previous formulation and propose an advanced formulation for participant selection in the PriCSS system to simultaneously achieve approximate social cost minimization, truthfulness, and differential location privacy. In what follows, we first outline some background knowledge to facilitate the presentation and understanding of our scheme. Then we present our advanced formulation with differential location privacy.

A. Background

Definition 1. An auction is truthful if and only if any bidder's (expected) utility of bidding its true valuation v_i is at least its (expected) utility of bidding any other value c_i [31],

$$u_i(v_i, c_{-i}) \geq u_i(c_i, c_{-i}). \quad (3)$$

Definition 2. A mechanism satisfies the voluntary participation condition if agents who bid truthfully never incur a net loss, i.e., $\text{profit}_i(v_i, (c_{-i}, v_i)) \geq 0$ for all agents i , true value v_i , and other agents' bids c_{-i} [32].

Clearly, the voluntary participation condition is a desired property of our scheme design.

Theorem 1. A decreasing output function admits a truthful payment scheme satisfying voluntary participation if and only if $\int_0^\infty x_i(c_{-i}, u) du \leq \infty$ for all i , c_{-i} . In this case, we can take the payments to be [32]

$$p_i(c_{-i}, c_i) = c_i x_i(c_{-i}, c_i) + \int_{c_i}^\infty x_i(c_{-i}, u) du \quad (4)$$

Differential privacy is a powerful tool to provide statistical guarantee on the privacy leakage induced by publishing outputs based on sensitive input data sets. The basic idea is that for two almost identical input data sets, the output of the mechanism are nearly identical. The formal definition of differential privacy is as follows [22].

Definition 3. A randomized function \mathcal{M} gives ϵ -differential privacy if for all data sets D_1 and D_2 differing on at most one element, and all $\mathcal{S} \subseteq \text{Range}(\mathcal{M})$,

$$\Pr[\mathcal{M}(D_1) \in \mathcal{S}] \leq \exp(\epsilon) \times \Pr[\mathcal{M}(D_2) \in \mathcal{S}]. \quad (5)$$

Approximate differential privacy relaxes on the strict requirement and allows a small additive term in the bound [33].

Definition 4. A randomized function \mathcal{M} gives δ -approximate ϵ -differential privacy if for all data sets D_1 and D_2 differing on at most one element, and all $\mathcal{S} \subseteq \text{Range}(\mathcal{M})$,

$$\Pr[\mathcal{M}(D_1) \in \mathcal{S}] \leq \exp(\epsilon) \times \Pr[\mathcal{M}(D_2) \in \mathcal{S}] + \delta. \quad (6)$$

The parameter δ ensures that although not all events can satisfy the strong guarantee as specified by Eq. (5), the alternation is only for very low probability cases. Hence, it is desired that ϵ and δ to be as close to 0 as possible.

The exponential mechanism is a powerful tool to facilitate mechanism design via differential privacy [23]. The query function defined here $q(A, r)$ maps a pair of input data set

A and candidate outcome r to a real valued “score,” with the understanding that the higher score is, the better performance the mechanism can achieve. Specifically, it is defined below.

$$\Pr[\varepsilon_q^\varepsilon(A) = r] \propto \exp(\varepsilon q(A, r)). \quad (7)$$

The exponential mechanism gives $2\varepsilon\Delta$ differential privacy, where Δ is the largest change in q by a single change of the input in A .

The following theorem suggests that the probability of a highly suboptimal output is exponentially low [34].

Theorem 2. *The exponential mechanism, when used to select an output $r \in R$, gives $2\varepsilon\Delta$ -differential privacy, letting R_{OPT} be the subset of R achieving $q(A, r) = \max_r q(A, r)$, ensures that*

$$\Pr[q(A, \varepsilon_q^\varepsilon(A)) < \max_r q(A, r) - \frac{\ln(|R|/|R_{\text{OPT}}|) - t}{\varepsilon}] \leq \exp(-t). \quad (8)$$

B. Differentially Private Participant Selection

Due to the NP-hardness of the basic problem, we propose an approximate algorithm, combined with the exponential algorithm, to achieve the desired approximate minimum social cost, low computation complexity, and differential privacy.

Algorithm 1 Participant Selection in PriCSS

Input: Universe set \mathcal{T} of sensing tasks, set $\mathcal{B} = \bigcup_{i \in \mathcal{N}} b_i$ of all submitted bids.

Output: Winner set \mathcal{W} , social cost c .

```

1: Initialization:  $\varepsilon' \leftarrow \frac{\varepsilon}{\Delta \cdot \varepsilon \ln(\varepsilon/\delta)}$ ,  $\mathcal{W} \leftarrow \emptyset$ ,  $c \leftarrow 0$ ,  $T_{\mathcal{W}} \leftarrow \emptyset$ ;

2: while  $|\mathcal{T} - T_{\mathcal{W}}| > 0$  do
3:   for all  $b_i$  in  $\mathcal{B}$  do
4:     if  $L_i \subseteq T_{\mathcal{W}}$  then
5:        $\mathcal{B} \leftarrow \mathcal{B} - \{b_i\}$ ;
6:     else
7:        $r(c_i) = \frac{c_i}{|(\mathcal{T} - T_{\mathcal{W}}) \cap L_i|}$ ;
8:     end if
9:   end for
10:  for all  $b_i$  in  $\mathcal{B}$  do
11:     $\Pr[\mathcal{W} \leftarrow \mathcal{W} \cup \{i\}] = \frac{\exp(-\varepsilon' \cdot r(c_i))}{\sum_{b_j \in \mathcal{B}} \exp(-\varepsilon' \cdot r(c_j))}$ ;
12:  end for
13:  Select  $b_i$  according to the computed probability distribution.
14:  if  $b_i$  is selected then
15:     $\mathcal{B} \leftarrow \mathcal{B} - \{b_i\}$ ;
16:     $\mathcal{W} \leftarrow \mathcal{W} \cup \{i\}$ ;
17:     $c = c + c_i$ ;
18:     $T_{\mathcal{W}} \leftarrow T_{\mathcal{W}} \cup L_i$ ;
19:  end if
20: end while
21: return  $\mathcal{W}$ ,  $c$ 

```

The objective of the PriCSS administrator is still to select a set of participants for bundled spectrum sensing tasks, and we refer to Section IV for the notation. We first define a

ranking metric to characterize the administrator’s preference for participants, which applies to participant $i \in [1, n]$:

$$r(c_i) = \frac{c_i}{|(\mathcal{T} - T_{\mathcal{W}}) \cap L_i|}, \quad (9)$$

where the set $T_{\mathcal{W}}$ denotes the set of subtasks included in the current winning bids, i.e., $T_{\mathcal{W}} = \bigcup_{i \in \mathcal{W}} L_i$.

The rationale of this definition is as follows. The administrator always tends to select the participant with the lowest claimed cost per subtask that has not yet been included in $T_{\mathcal{W}}$. In each iteration, each participant’s ranking preference is calculated. Then for any remaining participant i who has not been included in the winner list, we adopt the following quality score for the exponential mechanism,

$$q(c_i, x_i) = -r(c_i). \quad (10)$$

The “-” sign is placed to fit the exponential mechanism in our reverse auction model. It is clear that the smaller $r(c_i)$, the higher the quality score of participant i . This effect is preferred during the winner selection.

The details of the proposed allocation scheme is shown in **Algorithm 1**. According to the exponential mechanism, the probability of participant i being selected as a winner is

$$\Pr(x_i = 1) \propto \exp(-\varepsilon' r(c_i)), \quad (11)$$

where ε' is specified as $\frac{\varepsilon}{\Delta \cdot \varepsilon \ln(\varepsilon/\delta)}$. Δ is the maximum input difference for c_i , which equals $c_{\max} - c_{\min}$. ε and δ are parameters to balance the privacy leakage and efficiency (in terms of social cost minimization in our scenario). Line 11 in **Algorithm 1** can thus be derived considering all the unselected participants. It essentially normalizes the overall participants’ selection probability. Based on the selection probability for each remaining participant, participant i is selected as the winner in this iteration. We then remove his bid b_i from \mathcal{B} and include i in the winner set \mathcal{W} .

We resort to Theorem 1 for the truthful payment design. Each winner i is paid by the administrator with the amount

$$p_i(c_{-i}, c_i) = c_i x_i(c_{-i}, c_i) + \int_{c_i}^{c_{\max}} x_i(c_{-i}, u) du, \quad (12)$$

where $x_i(c_{-i}, c_i)$ represents the probability that participant i is selected to perform the sensing task bundle L_i when i ’s claimed cost is c_i and others’ claimed cost vector is c_{-i} .

VII. PERFORMANCE ANALYSIS

In this section, we prove how PriCSS achieves the desired design objectives: differential location privacy, approximate social cost minimization, and truthfulness.

A. Differential Location Privacy

Theorem 3. *For any $\delta \leq 1/2$, PriCSS preserves $((e - 1)\varepsilon' \Delta \ln(e\delta^{-1}), \delta)$ -differential location privacy.*

Proof: To facilitate the proof, we first define Q_i as the subtask set that participant i can still contribute to, i.e., $Q_i = (\mathcal{T} - T_{\mathcal{W}}) \cap L_i$. In two consecutive auction rounds, assume that there are two bidding vectors $\{c_1, c_2, \dots, c_l, \dots, c_n\}$ and

$\{c'_1, c'_2, \dots, c'_l, \dots, c'_n\}$ that differ by only one single element at the l th index. $c_i = c'_i$ for all $i \in [1, n]$ except $i = l$. Differential privacy suggests that with these two bidding vectors as input, the probability that the outputs of the mechanism, i.e., the winner sets \mathcal{W} and \mathcal{W}' , are approximately the same. The rationale of our proof is to obtain an exponential upper-bound for $\Pr[\mathcal{W} = \{w_1, w_2, \dots, w_p\}]/\Pr[\mathcal{W}' = \{w_1, w_2, \dots, w_p\}]$, where \mathcal{W} and \mathcal{W}' are the two ordered winner lists, i.e., w_i is always selected as a winner before w_j for any $j > i$. We give our formal proof below:

$$\begin{aligned} & \frac{\Pr[\mathcal{W} = \{w_1, w_2, \dots, w_p\}]}{\Pr[\mathcal{W}' = \{w_1, w_2, \dots, w_p\}]} \\ &= \prod_{i=1}^p \frac{\exp(-\varepsilon' \cdot c_i/|Q_i|)/\sum_{j \in \mathcal{N} \setminus \pi_i} \exp(-\varepsilon' \cdot c_j/|Q_j|)}{\exp(-\varepsilon' \cdot c'_i/|Q_i|)/\sum_{j \in \mathcal{N} \setminus \pi_i} \exp(-\varepsilon' \cdot c'_j/|Q_j|)} \\ &= \prod_{i=1}^p \frac{\exp(-\varepsilon' \cdot c_i/|Q_i|)}{\exp(-\varepsilon' \cdot c'_i/|Q_i|)} \cdot \prod_{i=1}^p \frac{\sum_{j \in \mathcal{N} \setminus \pi_i} \exp(-\varepsilon' \cdot c'_j/|Q_j|)}{\sum_{j \in \mathcal{N} \setminus \pi_i} \exp(-\varepsilon' \cdot c_j/|Q_j|)} \\ &= \exp(\varepsilon' \frac{c'_l - c_l}{|Q_l|}) \prod_{i=1}^p \frac{\sum_{j \in \mathcal{N} \setminus \pi_i} \exp(-\varepsilon' \cdot c'_j/|Q_j|)}{\sum_{j \in \mathcal{N} \setminus \pi_i} \exp(-\varepsilon' \cdot c_j/|Q_j|)}, \end{aligned} \quad (13)$$

where $\pi_1 = \emptyset$ and $\pi_i = \{w_1, w_2, \dots, w_{i-1}\} (i > 1)$. If $c_l < c'_l$, the second term is smaller than 1. Then

$$\frac{\Pr[\mathcal{W} = \{w_1, w_2, \dots, w_p\}]}{\Pr[\mathcal{W}' = \{w_1, w_2, \dots, w_p\}]} < \exp(\varepsilon' \Delta), \quad (14)$$

where Δ is the maximum difference of the bid values for the same set of task bundles.

If $c_l > c'_l$, the first term is smaller than 1. We denote $\alpha_j = c_j - c'_j$, then

$$\begin{aligned} & \frac{\Pr[\mathcal{W} = \{w_1, w_2, \dots, w_p\}]}{\Pr[\mathcal{W}' = \{w_1, w_2, \dots, w_p\}]} \\ &< \prod_{i=1}^p \frac{\sum_{j \in \mathcal{N} \setminus \pi_i} \exp(-\varepsilon' \cdot c'_j/|Q_j|)}{\sum_{j \in \mathcal{N} \setminus \pi_i} \exp(-\varepsilon' \cdot c_j/|Q_j|)} \\ &= \prod_{i=1}^p \frac{\sum_{j \in \mathcal{N} \setminus \pi_i} \exp(-\varepsilon' \cdot c'_j/|Q_j|)}{\sum_{j \in \mathcal{N} \setminus \pi_i} \exp(-\varepsilon' \cdot \alpha_j/|Q_j|) \exp(-\varepsilon' \cdot c'_j/|Q_j|)} \\ &= \prod_{i=1}^p \mathbb{E}_{j \in \mathcal{N} \setminus \pi_i} [\exp(\varepsilon' \cdot \alpha_j/|Q_j|)]. \end{aligned} \quad (15)$$

Note that for all $\eta \leq 1$, $e^\eta \leq 1 + (e-1)\eta$. Therefore, for all $\varepsilon' \leq 1/\Delta$,

$$\begin{aligned} & \frac{\Pr[\mathcal{W} = \{w_1, w_2, \dots, w_p\}]}{\Pr[\mathcal{W}' = \{w_1, w_2, \dots, w_p\}]} \\ &\leq \prod_{i=1}^p \mathbb{E}_{j \in \mathcal{N} \setminus \pi_i} (1 + (e-1) \cdot \varepsilon' \cdot \alpha_j) \\ &\leq \exp((e-1) \varepsilon' \sum_{i=1}^p \mathbb{E}_{j \in \mathcal{N} \setminus \pi_i} \alpha_j). \end{aligned} \quad (16)$$

So if $\sum_{i=1}^p \mathbb{E}_{j \in \mathcal{N} \setminus \pi_i} \alpha_j$ is upper-bounded, the theorem is established. Based on the proofs in [34], we have $\Pr(\sum_{i=1}^p \mathbb{E}_{j \in \mathcal{N} \setminus \pi_i} \alpha_j > \Delta \ln(e\delta^{-1})) \leq \delta$. ■

B. Approximate Social Cost Minimization

Theorem 4. With probability of at least $1 - 1/n^{\mathcal{O}(1)}$, PriCSS can assign spectrum sensing tasks to a set of winners with a social cost of at most $\gamma \text{OPT} + \mathcal{O}(\ln(n))$, where OPT denotes the optimal (minimum) social cost, and n is the number of participants.

Proof: Let \mathcal{W}_{OPT} denote the set of winners in the auction with the minimum social cost. We denote an arbitrary set of winners as \mathcal{W} and number the winners according to the order of being selected, i.e., $\mathcal{W} = \{w_1, w_2, \dots, w_l\}$.

For each $i \in \mathcal{W}$, we define a set \mathcal{W}_i , with the following constraints ($\forall j \in \mathcal{W}_i$):

- 1) $j \in \mathcal{W}_{\text{OPT}}$;
- 2) $Q_j \cap Q_i \neq \emptyset$;
- 3) $|Q_j - (Q_j \cap Q_i)| = 0$;
- 4) $Q_j \neq \emptyset$ before i is selected as one winner.

The above constraints suggest that in this arbitrary selection \mathcal{W} , the reason that a participant j is not listed is that there is a participant i with a conflicting task set with that of participant j , and i wins. Note that in Eq. (8), the q function corresponds to the inverse and unified cost in our scenario. Therefore, by taking $t = \mathcal{O}(\ln(n))$, we have

$$-\frac{c_i}{|Q_i|} \geq -\frac{c_j}{|Q_j|} - \mathcal{O}(\ln n) \quad (17)$$

with a probability of at least $1 - 1/n^{\mathcal{O}(1)}$.

Since $|Q_j|$ is upper bounded by a constant γ where $\gamma < n$ when n is large, we have

$$c_j \geq \frac{c_i}{|Q_i|} \cdot |Q_j| - \mathcal{O}(\ln n) \quad (18)$$

with a probability of at least $1 - 1/n^{\mathcal{O}(1)}$.

Summing all $j (j \in \mathcal{W}_i)$ together, we have

$$\begin{aligned} \sum_{j \in \mathcal{W}_i} c_j &\geq \left(\frac{c_i}{|Q_i|} - \mathcal{O}(\ln n)\right) \cdot \sum_{j \in \mathcal{W}_i} |Q_j| \\ &\geq \frac{c_i}{\gamma} - \mathcal{O}(\ln n). \end{aligned} \quad (19)$$

with a probability of at least $1 - 1/n^{\mathcal{O}(1)}$. The last step holds because $\sum_{j \in \mathcal{W}_i} |Q_j| \geq 1$.

Summing all $i \in \mathcal{W}$, we have

$$\begin{aligned} \sum_{j \in \mathcal{W}_{\text{OPT}}} c_j &= \sum_{i \in \mathcal{W}} \left(\sum_{j \in \mathcal{W}_i} c_j + \sum_{j \in \mathcal{W}_{\text{OPT}} \cap \mathcal{W}_i} c_j \right) \\ &\geq \sum_{i \in \mathcal{W}} \frac{c_i}{\gamma} - \mathcal{O}(\ln n). \end{aligned} \quad (20)$$

This concludes the proof. ■

C. Truthfulness

We finally prove that PriCSS is truthful. Based on Theorem 1, we need to show that the selection of PriCSS is monotone decreasing with an appropriate payment scheme.

Lemma 5. In PriCSS, for each participant i , the probability that i is assigned with the interested spectrum sensing task bundle is monotone decreasing with his claimed cost c_i .

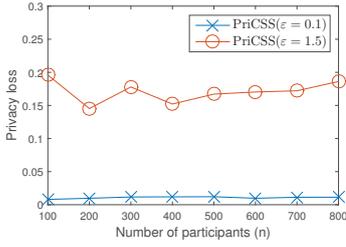


Fig. 2: Privacy loss.

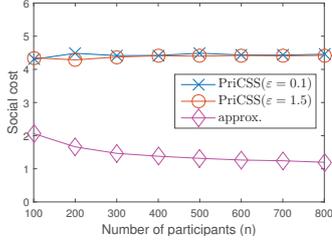


Fig. 3: Social cost for $K = 3$ sensing tasks.

Proof: Due to the randomized property of our scheme, we simply prove that the probability that i is assigned with the interested spectrum sensing task bundle is decreasing when his claimed cost c_i increases in each round of winner selection.

$$\begin{aligned}
 & \Pr(\mathcal{W} \leftarrow \mathcal{W} \cup \{i\}) \\
 &= \frac{\exp(-\varepsilon' \cdot r(c_i))}{\sum_{b_j \in \mathcal{B}} \exp(-\varepsilon' \cdot r(c_j))} \\
 &= \frac{\exp(-\varepsilon' \cdot r(c_i))}{\sum_{b_j \in \mathcal{B} \setminus \{c_i\}} \exp(-\varepsilon' \cdot r(c_j)) + \exp(-\varepsilon' \cdot r(c_i))} \quad (21) \\
 &= 1 - \frac{\sum_{c_j \in \mathcal{B} \setminus \{c_i\}} \exp(-\varepsilon' \cdot r(c_j))}{\sum_{b_j \in \mathcal{B} \setminus \{c_i\}} \exp(-\varepsilon' \cdot r(c_j)) + \exp(-\varepsilon' \cdot r(c_i))}
 \end{aligned}$$

In the above equation, if we increase c_i , $r(c_i)$ also increases. Then the exponential term of c_i decreases, causing the overall equation value to decrease. This indicates that if we increase c_i , the probability that \mathcal{W} includes i in every round decreases if i has not been included in previous rounds. ■

We thus have the following theorem established:

Theorem 6. *PriCSS is truthful.*

VIII. PERFORMANCE EVALUATION

In this section, we use simulations to evaluate whether PriCSS can achieve differential location privacy and approximate social cost minimization.

Our simulation setting is as follows. We simulate a square urban area of 1km by 1km. The PriCSS administrator issues sensing tasks in response to the queries of secondary users, each with a transmission radius of 300m. The base locations of PriCSS participants are uniformly distributed, and we vary the number of PriCSS participants from 100 to 800 in simulations. In our simulation, the preferred sensing locations are chosen beforehand according to the specific diversity requirement as discussed in Section III. To minimize the overall sensing cost, we want the subtask locations to be as far from each other

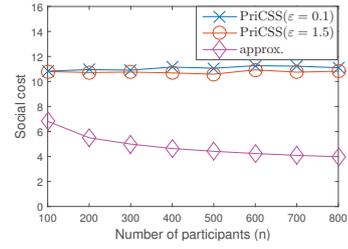


Fig. 4: Social cost for $K = 9$ sensing tasks.

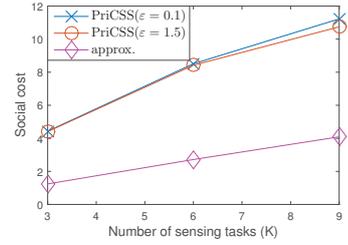


Fig. 5: Social cost for 700 participants.

as possible. We specify a minimum separation distance of 100m for the subtasks in each sensing task. The number of sensing locations (or subtasks) for each sensing task is fixed as 5 in our simulations. We vary the number of sensing tasks K in one round of auction from 3 to 9. Each sensing task is characterized by the locations of the corresponding secondary users, which are uniformly distributed within the region. We also set the modeling parameters $\eta = 100$ reward units and $\rho = 1$ unit per meter. In addition, we set the bidding cost range $[c_{\min}, c_{\max}]$ to $[100, 2000]$ and then normalize it to $[0, 1]$. The parameter γ is specified as 5. Note that other configurations of η and ρ lead to similar performance. We omit other cases here due to limited space. The privacy parameter ε is chosen as 0.1 or 1.5, and δ is set to 0.25. The simulations are done in MATLAB, and each result represents the average of 100 runs.

We use two metrics to evaluate PriCSS. The first is the privacy loss, defined according to **Definition 3**:

$$\epsilon = \max_{\mathcal{S}} \ln \frac{\Pr[\mathcal{M}(D_1) \in \mathcal{S}]}{\Pr[\mathcal{M}(D_2) \in \mathcal{S}]}, \quad (22)$$

where D_1 and D_2 correspond to two cost vectors for all the participants that differ by one element. Intuitively, the smaller ϵ , the less impact the change of single cost on the auction results, the better individual sensing cost privacy is protected, and the more location privacy each participant enjoys. The second metric is the social (or true sensing) cost of the winners, which is desired to be as low as possible. For the purpose of comparison, we also show the social cost induced using the approximation algorithm without privacy considerations introduced in Section IV. We first evaluate the location-privacy loss in PriCSS. As proved in Section VII, PriCSS preserves $((e-1)\varepsilon' \Delta \ln(e\delta^{-1}), \delta)$ -differential location privacy, where ε' is specified as $\frac{\varepsilon}{\Delta \cdot e \ln(e/\delta)}$. This is equivalent to achieving $(\frac{e-1}{e}\varepsilon, \delta)$ differential privacy. In the simulations, we set $\varepsilon = 0.1$ or 1.5 and $\delta = 0.25$. Fig. 2 shows the achievable privacy loss in PriCSS, which is obviously much

lower than the theoretical result. Specifically, when $\varepsilon = 0.1$, we can observe almost a constant privacy loss of 0.01, which is far lower than the theoretical value $\frac{\varepsilon-1}{10\varepsilon} \approx 0.06$. Similar conclusions can be drawn with $\varepsilon = 1.5$. This indicates that when there is any change of a single cost value for any participant, there is rarely any chance that the auction result can change. Since differential privacy mechanisms work for an arbitrary adversary, we can safely conclude that the attackers can no longer infer the participants' locations by performing the attacks in Section V or adopting other attack strategies.

We show the social cost incurred using PriCSS and the approximate algorithm (denoted as approx.) for three and nine sensing tasks in Fig. 3 and Fig. 4, respectively. For the approximate algorithm, we observe that as the number of participants increases, the social cost tends to decrease due to increased competition among participants. The trend of decrease, however, cannot be found with PriCSS for both $\varepsilon = 0.1$ and $\varepsilon = 1.5$ cases. We conjecture that with PriCSS in place, the advantage of cost-efficient participants who claim lower sensing costs in the hope of winning more is weakened by the increased number of participants. In other words, their ranking metrics play less significant roles when the number of participants increases. Still, we see that the social cost when $\varepsilon = 0.1$ is slightly worse than $\varepsilon = 1.5$. This is the expected trade-off between privacy and utility: the larger ε , the heavier weight on the ranking metric, and the lower the social cost. In Fig. 5, we also show the social cost for different numbers of sensing tasks when there are 700 participants. As expected, when the number of sensing tasks increases, the social cost also increases.

IX. CONCLUSIONS

In this paper, we present PriCSS, a novel framework for a spectrum database administrator to select spectrum-sensing participants in a differentially privacy-preserving manner. We provide detailed privacy and efficiency analysis of the scheme and evaluate the performance extensively. We demonstrate that PriCSS can simultaneously achieve the three design objectives: differential location privacy, approximate social cost minimization, and truthfulness.

X. ACKNOWLEDGEMENT

This work was supported in part by the US National Science Foundation under grants CNS-1514381, CNS-1421999, and CNS-1320906.

REFERENCES

- [1] T. Zhang, N. Leng, and S. Banerjee, "A vehicle-based measurement framework for enhancing whitespace spectrum databases," in *MobiCom'14*, Sept. 2014.
- [2] A. Nika, Z. Zhang, X. Zhou, B. Y. Zhao and H. Zheng, "Towards commoditized real-time spectrum monitoring," in *HotWireless'14*, Sept. 2014.
- [3] J. Sun, R. Zhang, X. Jin, and Y. Zhang, "SecureFind: Secure and privacy-preserving object finding via mobile crowdsourcing," *IEEE Trans. Wireless Communication*, Oct. 2015.
- [4] Z. Feng, Y. Zhu, Q. Zhang, L. M. Ni, and A. V. Vasilakos, "TRAC: truthful auction for location-aware collaborative sensing in mobile crowdsourcing," in *INFOCOM'14*, Apr. 2014.
- [5] S. Li, H. Zhu, Z. Gao, X. Guan, K. Xing, and X. Shen, "Location privacy preservation in collaborative spectrum sensing," in *INFOCOM'12*, Mar. 2012.
- [6] Z. Gao, H. Zhu, S. Li, and S. Du, "Security and privacy of collaborative spectrum sensing in cognitive radio networks," *IEEE Trans. Wireless Communications*, Dec. 2012.
- [7] W. Wang, and Q. Zhang, "Privacy-preserving collaborative spectrum sensing with multiple service providers," *IEEE Trans. Wireless Communications*, Oct. 2014.
- [8] H. To, G. Ghinita, and C. Shahabi, "A framework for protecting worker location privacy in spatial crowdsourcing," *VLDB'14*, Sept. 2014.
- [9] R. Zhang, J. Zhang, Y. Zhang, and C. Zhang, "Secure crowdsourcing-based cooperative spectrum sensing," in *INFOCOM'13*, Apr. 2013.
- [10] X. Wang, W. Cheng, P. Mohapatra, and T. Abdelzaher, "ARTSense: Anonymous reputation and trust in participatory sensing," in *INFOCOM'13*, Apr. 2013.
- [11] R. Chen, J. Park, and K. Bian, "Robust distributed spectrum sensing in cognitive radio networks," in *INFOCOM'08*, Apr. 2008.
- [12] A. Min, K. Shin, and X. Hu, "Attack-tolerant distributed sensing for dynamic spectrum access networks," in *ICNP'09*, Oct. 2009.
- [13] H. Li, and Z. Han, "Catch me if you can: an abnormality detection approach for collaborative spectrum sensing in cognitive radio networks," *IEEE Trans. Wireless Communications*, Nov. 2010.
- [14] S. Li, H. Zhu, Z. Gao, X. Guan, and K. Xing, "YouSense: mitigating entropy selfishness in distributed collaborative spectrum sensing," in *INFOCOM'13*, Apr. 2013.
- [15] Y. Liu, P. Ning, and H. Dai, "Authenticating primary users' signals in cognitive radio networks via integrated cryptographic and wireless link signatures," in *S&P'10*, May. 2010.
- [16] V. Kumar, J. Park, and K. Bian, "Blind transmitter authentication for spectrum security and enforcement," in *CCS'14*, Nov. 2014.
- [17] X. Jin, J. Sun, R. Zhang, Y. Zhang, and C. Zhang, "SpecGuard: Spectrum misuse detection in dynamic spectrum access systems," in *INFOCOM'15*, Apr. 2015.
- [18] X. Jin, J. Sun, R. Zhang, and Y. Zhang, "SafeDSA: Safeguard dynamic spectrum access against fake secondary users," in *CCS'15*, Oct. 2015.
- [19] X. Zhang, G. Xue, R. Yu, D. Yang, and J. Tang, "Truthful incentive mechanisms for crowdsourcing," in *INFOCOM'15*, Apr. 2015.
- [20] D. Yang, G. Xue, X. Fang, and J. Tang, "Crowdsourcing to smartphones: incentive mechanism design for mobile sensing," in *MobiCom'12*, Aug. 2012.
- [21] D. Zhao, X. Li, and H. Ma, "How to crowdsource tasks truthfully without sacrificing utility: online incentive mechanisms with budget constraint," in *INFOCOM'14*, Apr. 2014.
- [22] C. Dwork, "Differential privacy," in *ICALP'06*, July 2006.
- [23] F. McSherry, K. Talwar, "Mechanism design via differential privacy," in *FOCS'07*, Oct. 2007.
- [24] Z. Huang, S. Kannan, "The exponential mechanism for social welfare: private, truthful, and nearly optimal," in *FOCS'12*, Oct. 2012.
- [25] R. Zhu, Z. Li, F. Wu, K. G. Shin, and G. Chen, "Differentially private spectrum auction with approximate revenue maximization," in *MobiHoc'14*, Aug. 2014.
- [26] R. Zhu, K. G. Shin, "Differentially private and strategy-proof spectrum auction with approximate revenue maximization", in *INFOCOM'15*, Apr. 2015.
- [27] J. Sun, R. Zhang, J. Zhang, and Y. Zhang, "PriStream: Privacy-preserving distributed stream monitoring of thresholded percentile statistics," in *INFOCOM'16*, Apr. 2016.
- [28] Y. Selen, H. Tullberg, and H. Kronander, "Sensor selection for cooperative spectrum sensing," in *DySPAN'08*, Oct. 2008.
- [29] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, "Algorithmic game theory," Cambridge University Press, 2007.
- [30] V. Vazirani, "Approximation algorithms," Springer-Verlag, ISBN 3-540-65367-8.
- [31] J. Jia, Q. Zhang, Q. Zhang, and M. Liu, "Revenue Generation for Truthful Spectrum Auction in Dynamic Spectrum Access," in *MobiHoc'09*, May 2009.
- [32] A. Archer, É. Tardos, "Truthful mechanisms for one-parameter agents," in *FOCS'01*, Oct. 2001.
- [33] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, "Our data, ourselves: privacy via distributed noise generation," in *EUROCRYPT'06*, May 2006.
- [34] A. Gupta, K. Ligett, F. McSherry, A. Roth, and K. Talwar, "Differentially private combinatorial optimization," in *SODA'10*, Nov. 2009.