ALETHEIA: Robust Large-Scale Spectrum Auctions against False-name Bids

Qinhui Wang [†♯], Baoliu Ye[†*], Bin Tang[†], Tianyin Xu[‡], Song Guo[¶], Sanglu Lu[†], and Weihua Zhuang[§]
[†]National Key Laboratory for Novel Software Technology, Nanjing University, China
[♯]Department of Military Management, Army Command College, China
[‡]Department of Computer Science and Engineering, U.C. San Diego, USA
[¶]School of Computer Science and Engineering, The University of Aizu, Japan
[§]Department of Electrical and Computer Engineering, University of Waterloo, Canada
qhwang@dislab.nju.edu.cn, {yebl,tb, sanglu}@nju.edu.cn, tixu@cs.ucsd.edu, sguo@u-aizu.ac.jp, wzhuang@bbcr.uwaterloo.ca

ABSTRACT

Auction is a promising approach for dynamic spectrum access in Cognitive Radio Networks. Existing auction mechanisms are mainly proposed to be *strategy-proof* to stimulate bidders to reveal their valuations of spectrum truthfully. However, they would suffer significantly from a new cheating pattern named *false-name bids*, where a bidder can manipulate the auction by submitting bids under multiple fictitious names. We show such false-name bid cheating is easy to make but hard to be detected in dynamic spectrum auctions. To resolve this issue, we propose ALETHEIA, a novel *flexible, false-name-proof* auction framework for large-scale dynamic spectrum access. ALETHEIA has the following important features: (1) it not only guarantees strategy-proofness but also resists false-name bids, (2) it enables spectrum reuse across a large number of bidders, (3) it provides the bidders the flexibility of diverse demand formats, and (4) it incurs low computational overhead. Simulation results show that ALETHEIA achieves both high spectrum redistribution efficiency and auction efficiency.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design

Keywords

Spectrum auctions, false-name-proofness, cognitive radio networks

1. INTRODUCTION

1.1 Motivation

Radio spectrum is a critical but scarce resource for wireless communication. With the rapid growth of wireless services and devices, the limited spectrum is draining away, while most licensed spectrum bands (e.g., TV channels) are under-utilized as reported in [2, 14]. To address this dilemma, dynamic spectrum access (DSA) has been proposed based on the advance of cognitive radio (CR) techniques [2], where primary users with licenses can gain financial benefits by leasing their idle spectrum to secondary users without licenses to access, i.e., a win-win situation.

Auctions are the well recognized and *de facto* mechanisms for redistributing spectrum in DSA that can achieve both fairness and allocation efficiency [10]. In auction-based DSA, the spectrum is divided into multiple channels, then the secondary users (referred to as bidders) submit their bids for channels based on the valuation of their short-term local usage, and finally the auction mechanism performed by the primary user (referred to as auctioneer) determines winners as well as their channel allocation and payment. In order to achieve fairness and maximize spectrum utilization, existing spectrum auction designs (e.g., [3, 9, 20, 24]) mainly aim to i) exploit *spectrum reusability*, i.e., a channel can be allocated to multiple bidders as long as these bidders do not interfere with each other, and ii) guarantee *strategy-proofness* (a.k.a. *truthfulness*), i.e., no bidder can improve its own utility by cheating with false valuation, so that each bidder is encouraged to tell the auctioneer its true valuation.

However, the auction manipulation by the bidders could go beyond the cheating with false valuations. As demonstrated in recent studies [5, 7, 16], it is fairly easy for a CR user to generate multiple "names" identified by service-set identifiers (SSIDs). For example, the Atheros chipset supports up to 64 identifiers for one physical device [1]. This can lead to a new cheating pattern named *false-name bids*, i.e., a bidder can submit bids using multiple fictitious names to manipulate the auction results. Unfortunately, due to the

*B. Ye is the corresponding author.
[†]This work was partially supported by NSFC Grants (No. 61170069, 61373014, 91218302 and 61321491); Science and Technology Pillar Program (Industry) of Jiangsu Province Grant (No. BE2013116); Collaborative Innovation Center of Novel Software Technology and Industrialization; EU FP7 IRSES MobileCloud Project Grant (No. 612212).

MobiHoc'15, June 22–25, 2015, Hangzhou, China.

http://dx.doi.org/10.1145/2746285.2746307.

open, mobile and ubiquitous nature of CR users, it is prohibitively difficult to detect false identifiers, even using the state-of-the-art authentication methods [5,7,16]. A series of questions then naturally arise. What are the consequences of false-name bids for existing spectrum auction mechanisms? Can we design efficient auction mechanisms that are *false-name-proof*, so that each bidder cannot improve its utility by using false-name bids?

1.2 Our results

To answer the first question, we conduct a broad and deep investigation on several typical strategy-proof spectrum auction mechanisms. For example, we use simulation experiments to examine the effect of false-name cheating to VERITAS [24], a state-of-the-art strategy-proof spectrum auction design. The results show that even a simple pattern of false-name cheating can substantially improve the utilities of cheating bidders and reduce the revenue of the auctioneer by 40%. This highlights the importance of providing false-name-proofness in spectrum auction designs. While prior studies have enforced false-name-proofness in auctions for traditional goods (e.g., books and paintings) [11,17], we show that existing false-name-proof designs fail when spectrum reusability is embraced. Moreover, it usually results in a low spectrum efficiency.

The above pessimistic results lead us to develop a completely new spectrum auction to resist false-name-bids robustly. In this paper, we propose ALETHEIA, the *first* false-name-proof seal-bid spectrum auction framework for the emerging large-scale DSAs. It is implemented in a price oriented fashion, where the prices of bidders are computed first while winners are then determined independently based on these prices and get corresponding channels in a sequential manner. In this way, it is then ensured that the price of buying a bundle of channels is no larger than the sum of prices for buying these channels separately using multiple fictitious names. In summary, ALETHEIA provides the following advantages:

- As the primary feature, both strategy-proofness and false-name-proofness are guaranteed.
- It exploits the spectrum reusability substantially, to significantly improve spectrum utilization.
- It provides the bidders the flexibility of diverse demand formats, i.e., either an exact number of or a range of channels to be asked.
- It incurs a low computational overhead, making it more applicable to real-time spectrum auction in large-scale DSA. For instance, when each bidder requests an exact number of channels, it has the same time complexity as VERITAS at order of $O(N \log N + NKE)$, where N, K, E are the number of bidders, channels, and pairs of conflicting bidders (i.e., they cannot access the same channel simultaneously), respectively.
- It achieves both high spectrum redistribution efficiency and auction efficiency, as demonstrated by extensive simulation results.

2. PRELIMINARIES

2.1 System Model

We consider a cognitive radio network where spectrum is divided into K identical channels, denoted as $\mathcal{K} = \{1, \ldots, K\}$ and to be auctioned to N secondary users, denoted as $\mathcal{N} = \{1, \ldots, N\}$. We assume that each bidder $i \in \mathcal{N}$ requests $d_i (0 < d_i \leq K)$ channels, and it has a *valuation* function $v_i : \mathcal{K} \to \mathbb{R}^+$ which calculates the true valuation (a non-negative value) for the requested channels. Each bidder submits its bid valuation b_i to the auctioneer, with a per-channel bid valuation t_i ($t_i = b_i/d_i$). Note that the bid valuation b_i does not have to be equal to the true valuation $v_i(d_i)$ if manipulation on b_i can make profit.

After collecting all bids and requests that are submitted by all bidders simultaneously, the auctioneer determines the winners from the bidders based on the predefined allocation rules. It then charges each bidder i with a payment denoted as $p_i(b_i, \mathbf{b}_{-i})$, where $\mathbf{b}_{-i}$ denotes the vector including all the bids from b_1 to b_N except b_i. Note that no payment applies to loser i, i.e., $p_i(b_i, \mathbf{b}_{-i}) = 0$. The *utility* of bidder i, denoted by $u_i(b_i, \mathbf{b}_{-i})$, is then defined as the difference between valuation and payment, *i.e.*, $u_i(b_i, \mathbf{b}_{-i}) = v_i(d_i) - p_i(b_i, \mathbf{b}_{-i})$. We simplify the notations $v_i(d_i)$, $p_i(b_i, \mathbf{b}_{-i})$, and $u_i(b_i, \mathbf{b}_{-i})$ as v_i, p_i, and u_i respectively, if no confusion is incurred.

Same as [8,19,24,25], we use the conflict graph $G(\mathcal{N}, \mathcal{E})$ [8] to capture the interference among bidders, where $\mathcal{N}$ is the set of bidders and $\mathcal{E}$ is the collection of all edges. An edge $(i, j) \in \mathcal{E}$ iff bidders i and j conflict with each other when they use a same channel simultaneously. Therefore, a channel can be allocated to multiple bidders as long as no edge exists between any pair of these bidders, which is referred to as *spatial reusability.* Without loss of generality, we assume that the conflict graph is a connected graph. Otherwise, we can perform the auction in each connected component.

When a bidder uses multiple names to submit bids, the auctioneer takes them as if they are from a group of bidders, called *virtual bidders* in this paper. These virtual bidders have the same interference condition of the real bidder, *i.e.,* the real bidder's conflicting neighbors still conflict with these virtual bidders. This is because when a bidder generates multiple names, its geographic information does not change. Moreover, to obtain different channels for real use, these bidders conflict with each other, *i.e.,* the subgraph consisting of these virtual bidders is a clique.

2.2 Design Targets

The basic target of our auction design is to satisfy *allocation feasibility.*

Definition 1 (Allocation Feasibility). *An allocation* $\} = (\mathcal{A}_1, \ldots, \mathcal{A}_N)$, *where* $\mathcal{A}_i \subset \mathcal{K}$ *denotes the set of channels allocated to bidder i, is feasible if* $\mathcal{A}_i \cap \mathcal{A}_j = \emptyset$ *holds for any* $(i, j) \in \mathcal{E}$.

More importantly, we need to ensure strategy-proofness and false-name-proofness that are the primary design targets of the auction mechanisms proposed in this paper.

Definition 2 (Strategy-proofness [13, 15]). *An auction mechanism is strategy-proof (or truthful) if for any bidder i and* $\mathbf{b}_{-i}$, $u_i(v_i, \mathbf{b}_{-i}) \geq u_i(b_i, \mathbf{b}_{-i})$ *holds for any* $b_i \neq v_i$.

Definition 3 (False-name-proofness [18, 22]). *An auction mechanism is false-name-proof if for any bidder i using m false identifiers $i_1, \dots, i_m$ to participate the auction and any $\mathbf{b}_{-i}$,*

$$u_i(v_i, \mathbf{b}_{-i}) \geq \sum_{j=1}^{m} u_{i_j}(b_{i_j}, \mathbf{b}_{-i} \cup I_{-j}^m)$$

where $I_{-j}^m = \{b_{i_l} : l \in \{1, \dots, m\}, l \neq j\}$.

Generally speaking, strategy-proofness prohibits improved utility from cheating on bid valuation, while false-name-proofness not only prohibits such a cheating but also discourages bidders to submit false-name bids. Note that false-name-proofness generalizes the concept of strategy-proofness (the case when $m = 1$). In other words, false-name-proofness is a sufficient but in general not a necessary condition of strategy-proofness.

3. FALSE-NAME BID CHEATING IN SPECTRUM AUCTIONS

In this section, we study the formation and the impact of false-name cheating in the emerging dynamic spectrum auctions. We demonstrate that *a simple form of false-name cheating can effectively raise the cheating bidders' utility and thus degrade the auctioneer's revenue (i.e., the sum of payments of the winners).*

3.1 Methodology

We start by demonstrating that it could be fairly easy to design a false-name cheating pattern in existing strategy-proof spectrum auctions that can raise bidder's utility effectively. Then conduct simulation-based experiments to examine how this pattern of cheating impacts the auction outcomes.

Since a cheating pattern is specific to the auction design, we select VERITAS [24], a most famous strategy-proof spectrum auction mechanism, for our case study. Similar works are also done to some other mechanisms, e.g., SMALL [20], which lead to very similar results as in this paper, so we omit the details for saving space. For the sake of completeness, we briefly introduce how VERITAS works. In VERITAS, bidders are considered in the descending order of per-channel bid, and if there are enough available channels for a bidder, the bidder will get the requested number of channels with smallest indices. A channel is available to a bidder if it has not been assigned to any other neighboring bidders. To charge a winner i, VERITAS finds its critical neighbor, the first rejected neighboring bidder j who would win if i do not participate the auction, and then charges i with $d_i \cdot t_j$. If no such a neighbor exists, then charge 0.

3.2 A Cheating Pattern

We first introduce a very simple false-name cheating pattern in VERITAS named Real-Fake (RF) cheating. It works as follows. A bidder i uses two identifiers, Real and Fake, to bid. Then Real and Fake bid x ($0 < x < d_i$) and $d_i - x$ channels, respectively, with the same per-channel bid t_i. If bidder i can win using a single identifier, it is straightforward to check that it can also win using RF cheating, and achieve the same or higher utility. Figure 1 shows such an example where higher utility is obtained via RF cheating, where 3 bidders (A,B,C) competing for 3 channels. When bidding via a single name, bidder A obtains a utility of 3 with bidder B as its critical neighbor. However, when bidder A uses two identifiers (A_1 and A_2) to bid, bidder A achieves an improved utility $u_A = u_{A_1} + u_{A_2}$=2+3=5 since A_1's critical neighbor is B while A_2 does not have a critical neighbor.

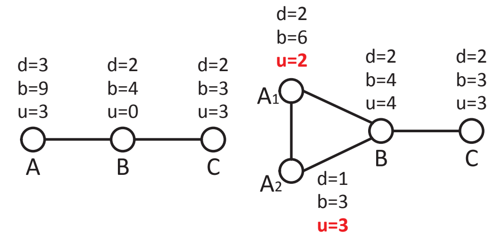

Figure 1: An example of RF cheating in VERITAS. When bidder A uses two names (A_1 and A_2) to bid, it increases its utility by getting the same number of channels but paying less.

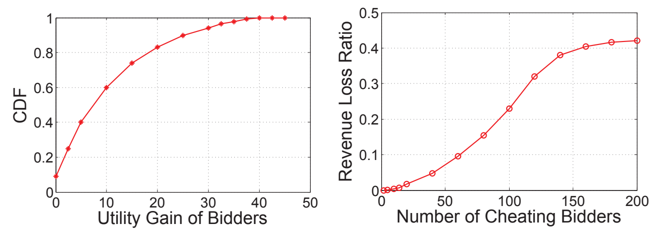

(a) Utility Gain (b) Revenue loss

Figure 2: RF cheating is effective in large-scale spectrum auctions like VERITAS [24]. (a) It can effectively improve bidder utility, giving bidders incentive to cheat. (b) It can significantly impair the revenue when a number of bidders choose to cheat.

3.3 Impacts of False-name Cheating

We further conduct simulations to examine how RF cheating impairs the auction outcomes. We simulate 1000 bidders competing 10 channels in the auction. These bidders are positioned in a 100×100 square and the interference range is set as 1, where each bidder has about 3 conflict neighbors on average, and thus a high degree of spectrum reuse could be exploited. The request and per-channel bid of each bidder are randomly chosen from integers in $[1, 6]$ and $[1, 10]$, respectively. For RF cheating, the value x is also uniformly chosen at random from all possible values.

We run the experiment 100 times. Figure 2(a) plots the utility gain of each RF cheating bidder. We observe that over 80% cheating bidders can gain profit by submitting false-name bids. Figure 2(b) plots the revenue loss ratio where the number of cheating bidders varies. From this figure, we can see that the revenue loss depends heavily on the number of cheating bidders. When the number of cheating bidders is small, the revenue loss ratio is low. This is because the effect of each RF cheating has a local feature which is inherited from the pricing method of VERITAS. When the number of RF cheating grows, the revenue decreases quickly, and the loss ratio can achieve 35% -40% when more than 150 bidders perform RF cheating.

The above results confirm that the unique requirement of spectrum reuse and the resulting local competition make traditional mechanisms vulnerable to false-name bids. Even simple RF cheating can significantly impair the auction revenue and fairness. This motivates us to design new mechanisms that are robust to the false-name cheating.

4. CHALLENGES

In this section, we discuss the challenges to design false-name-proof spectrum auction mechanisms. We first analyze the classic false-name-proof auctions and articulate their limitations when spectrum reusability is exploited. Then we introduce a simple false-name-proof design with fairly low spectrum reuse, which serves as a baseline for performance comparison.

4.1 False-name-proofness vs. Spatial Reusability

We first study two classic false-name-proof designs, GAL [17] and IR [11] for conventional goods auctions and show that they cannot even guarantee truthfulness and thus lose false-name-proofness when spatial reuse is enabled in spectrum auctions. One basic rule for extension here is to guarantee that the extended mechanism is equivalent to the original one if all bidders conflict with each other, *i.e.,* the conflict graph is a complete graph.

4.1.1 GAL Auction Design

GAL [17] is a false-name-proof multi-unit auction design using greedy allocation. It can be extended for spectrum auctions, with the following allocation and pricing methods.

Allocations:

1. Sort bidders in a descending order by the per-channel bid and set each bidder's available channel set as $\mathcal{K}$.

2. Check these sorted bidders one by one. For each bidder i, if there are at least d_i channels in its available-channel set, it wins and is allocated with the first d_i channels with the lowest indices, each of which is then removed from the available-channel sets of the neighbors of i. Otherwise, bidder i loses.

Pricing: For each winner i, find the first rejected conflicting neighbor j with a rank behind i. If such neighbor exists, then charge $d_i \cdot t_j$. Otherwise, charge 0.

We prove that this auction is not false-name-proof using a counter example, in which five bidders (A, B, ąd' ąd' ąd' , E) compete for two channels, as show in Figure 3. We find out that bidder E will lose with a utility of 0 under truthful bidding, while increasing its utility to 1 if it cheats by raising its bid to 4. In the latter case, bidder E will obtain a channel with a charge of 0. Hence, this mechanism is not strategy-proof and thus not false-name-proof.

4.1.2 IR Auction Design

Iterative Reducing (IR) [11] is another false-name-proof multi-unit mechanism where bidders with largest requests will be considered firstly. We extend IR by following this idea while considering spatial reuse, leading to the following auction design.

Allocation:

1. Group bidders by their requests d_i and sort these groups in a descending order by d_i. In each group, all bidders are sorted in a descending order by per-channel bid t_i. Initial each bidder's available channel set as $\mathcal{K}$.

2. Sequentially check the ordered groups. For a group, channels are allocated to the bidders sequentially until all bidders in the group win or a bidder loses. In the former case, we continue to check the next group. In the latter case, all the remaining bidders in the group lose and the allocation terminates.

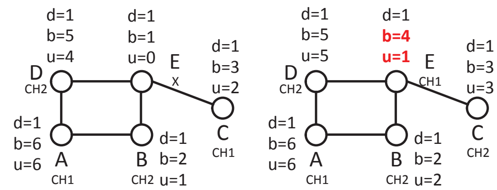

Figure 3: An illustrative example shows extended GAL is not false-name-proof, because bidder E can improve its utility by raising its bid.

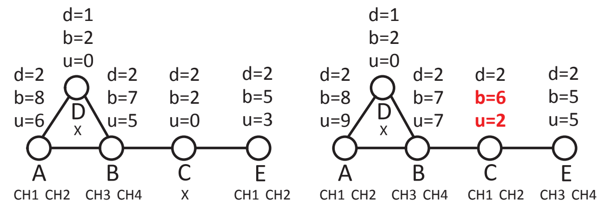

Figure 4: An illustrative example shows extended IR is not false-name-proof, where bidder C can improve its utility by raising its bid.

Pricing: If all bidders in a same group win, then these bidders are charged 0. Otherwise, each winner i in this group is charged by $d_i * t_j$, where j is the first rejected bidder in the group.

Again, we prove the above mechanism is not false-name-proof, using a counter example given in Fig. 4, where 5 bidders $A, B, \cdots, E$ are competing 4 channels. By the grouping method, bidders A, B, C and E are in group g_1, and bidder D constitutes group g_2. Consider g_1. Bidder C loses and its utility is 0. However, if C cheats by bidding 6, all bidders in g_1 win and thus are charged 0. Therefore, C improves its utility to 2, breaking the strategy-proofness and thus the false-name-proofness.

4.2 A False-name-proof Spectrum Auction

The spectrum utilization would be significantly impaired when existing false-name-proof designs are directly applied to spectrum auctions. The mechanism works as follows.

- Plane Division and Coloring. Divide the plane into boxes with length of the maximal interference radius, and proceed to uniformly color these boxes using 4 colors. The coloring guarantees: i) each pair of bidders in the same box interfere with each other, ii) bidders in different co-colored boxes do not interfere with each other. See Figure 5 for an illustration.

- Allocation in Each Box. Split all K channels into 4 subsets, each with $K/4$ channels. Allocate these sub-

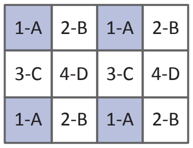

Figure 5: The plane is divided into boxes and colored using 4 colors (indexed from 1 to 4). The channels are divided into 4 subsets (indexed from A to D). Allocation rule is one color relates one subset.

sets to the boxes, one color relating to an independent subset. In each box, apply the GAL mechanism.

Since the false-name-proofness is guaranteed within each box by the GAL, when a bidder in a box submits false-name bids, the created virtual bidders are all in the same box and thus cannot improve its utility by the GAL. Therefore, the above mechanism is false-name-proof. Nevertheless, the static partition of spectrum among the boxes could result in a significant spectrum degradation of a factor up to 4.

Algorithm 1: ALETHEIA-Prices(B, i)

1 $p_i \leftarrow 0$; $\mathcal{A}_{\mathbb{N}(i)} \leftarrow \emptyset$;
2 **for** $i \in \mathcal{N}$ **do**
3     $Avai(i) \leftarrow \mathcal{K}$;
4 $B^{'} \leftarrow B \backslash \{i\}$;
5 **while** $B^{'} \neq \phi$ **do**
6     $j \leftarrow Top(B^{'})$;
7   **if** $|Avai(j)| \geq d_j$ **then**
8   Let $\mathcal{C}$ be the set of d_j channels with lowest indices in $Avai(j)$;
9   **if** $j \notin \mathbb{N}(i)$ || $(j \in \mathbb{N}(i) \&\& |\mathcal{C} \cup \mathcal{A}_{\mathbb{N}(i)}| + d_i \leq K)$ **then**
10             $\mathcal{A}_i \leftarrow C$;
11   **for** $k \in \mathbb{N}(j)$ **do**
12                 $Avai(k) \leftarrow Avai(k) - \mathcal{C}$;
13     $B^{'} \leftarrow B^{'} \backslash \{j\}$;
14 **for** $j \in \mathbb{N}(i)$ **do**
15   **if** $A_j = \emptyset$ **then**
16         $temp \leftarrow d_i * b_j / d_j$;
17         $p_i \leftarrow max(p_i, temp)$;
18 **return** p_i;

5. ALETHEIA

Motivated by the findings in previous Sections, we design ALETHEIA, a novel false-name-proof mechanism for large-scale dynamic spectrum auctions. In this section, we introduce the basic design of ALETHEIA, where bidders are assumed to be *single-minded* [11, 17, 24], i.e., a bidder i that requests d_i channels only accepts either all d_i channels or nothing. The extension of ALETHEIA to support range-based requests where bidder i accepts any number of channels between 0 and d_i is deferred to the next section.

5.1 Auction Design

ALETHEIA consists of three steps: (1) order the bidders into a bidder list (§5.1.1), (2) compute the price for each bidder (§5.1.2), and (3) determine the winners according to the finely computed prices (§5.1.3). One example is given in §5.1.4 to show how ALETHEIA works.

5.1.1 Bidder Ordering

A fixed order of bidders is critical to satisfying the interference constraints and makes it easier to design an bid-independent pricing rule. To maximize spectrum reusability, we present a novel ordering algorithm based on the Breadth-First-Search (BFS) procedure.

The ordered bidder list B is built via constructing a tree described as follows. (1) The bidder with the largest per-channel bid is selected as the root node, and all its conflicting neighbors become its child nodes. (2) For each leaf node, from the one with largest per-channel bid to the lowest one, we add its unincluded conflicting neighbors as its child nodes. The second step iteratively proceeds until all bidders are included in the tree. Note that a bidder's child nodes are ordered in a descending order of per-channel bid from the left to the right.

Finally, list B is obtained by walking through the tree layer by layer from the root node and from left to right on each layer. In particular, whenever we visit a new node, we append it at the tail of a FIFO list. The resulting list is B, after we search the whole tree.

5.1.2 Pricing Scheme

With the sorted bidder list B, we use the pricing algorithm (Algorithm 1) to compute the price for each bidder. The price for each bidder i is its request d_i multiplied by the per-channel bid of its *critical bidder*, denoted as $c(i)$. Generally speaking, a bidder i will win when bidding higher than its critical bidder $c(i)$ (*i.e.*, $t_i > t_{c(i)}$), and it will lose when bidding lower than that.

[Find Critical Bidder] To find the critical bidder for i, we first suppose that d_i channels have been allocated to bidder i, to reserve d_i channels for bidder i. It proceeds to allocate channels to other bidders sequentially, and find its all neighboring bidders that cannot be allocated. The losing neighboring bidder with the largest per-channel bid is i's critical bidder. If all of its neighbors win, then the critical bidder is a *null* bidder, whose per-channel bid equals 0.

The pricing algorithm, described in Algorithm 1 where $\mathcal{A}_{\mathbb{N}(i)}$ represents the set of channels that have been allocated to i's neighbors, *i.e.*, $\mathcal{A}_{\mathbb{N}(i)} = \cup_{j \in \mathbb{N}(i)} \mathcal{A}_j$, sequentially examine each bidder j ($j \neq i$). When there are enough channels for j (line 7), there are two cases to be considered (line 9). Case 1: If j does not belong to $\mathbb{N}(i)$, we assign the d_j channels with the **lowest indices** in $Avai(j)$ to bidder j. Case 2: If j conflicts with i, since we have to reserve d_i channels for i, we thus need to check whether j's allocation would cause i's request cannot be satisfied. In other words, we need to ensure the equation $|\mathcal{C} \cup \mathcal{A}_{\mathbb{N}(i)}| + d_i \leq K$ holds when allocating $\mathcal{C}$ to bidder j, where $\mathcal{C}$ denotes the channels pre-allocated to bidder j. After all bidders except i have been considered, we find the losing neighbor j with largest per-channel bid, and compute the price as $d_i * t_j$ (line 14-17).

5.1.3 Allocation Rule

Channel allocation is made by applying Algorithm 2 that sequentially checks all bidders. For each bidder i, the algorithm examines whether its bid valuation is greater than its computed price. If so, we assign d_i channels **with lowest indices** in its available channel set $Avai(i)$ to bidder i.

Algorithm 2: ALETHEIA-Allocation(B)

```
for i ∈ 𝒩 do
    Avai(i) ← 𝒦;
while B ≠ φ do
    i ← Top(B);
    if b_i > p_i then
        Set 𝒜_i as the set of channels with lowest indices
        in Avai(i);
        Assign 𝒜_i to bidder i;
        for j ∈ ℕ(i) do
            Avai(j) ← Avai(j) − 𝒜_i;
    B ← B\{i};
```

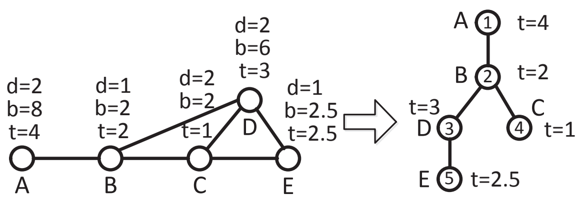

Figure 6: An illustrative example of ALETHEIA, where 3 channels (CH1, CH2 and CH3) are auctioned: (left) the conflict graph; (right) the corresponding BFS tree. The index of each bidder in list B is given in the node circle.

Otherwise, bidder i loses with no charge. In this way, we ensure that no bidders would be charged more than its bid valuation, to incentivize bidders to participate the auction.

5.1.4 All Together in One Example

Consider an example shown in Figure 6(left), where 5 bidders compete for 3 channels. The ordered bidder list is shown in Figure 6(right). We use bidder D as an example to find its critical bidder. We sequentially allocate channels to other bidders by reserving 2 channels for bidder D. Then bidder C and bidder E lose. Since E has higher per-channel bid, it becomes D's critical bidder. Similarly, we can find that *null*, C, D and C are the critical bidders of A, B, C and E, respectively. Therefore, we know bidder C loses and other bidders win by the allocation rule. Bidders A and D win CH1 and CH2, and bidders B and E win CH3.

5.2 Properties

Now we introduce the key properties of ALETHEIA. The proofs are deferred to the next subsection. The basic property regarding the correctness of ALETHEIA is that it can always produce a feasible allocation.

THEOREM 1. *ALETHEIA satisfies allocation feasibility, i.e., for any bidder $i \in \mathcal{N}$ with $b_i > p_i$, its request will be satisfied without causing conflicts.*

The key feature of ALETHEIA is its false-name-proofness, given in the following theorem.

THEOREM 2. *ALETHEIA is false-name-proof.*

Another property of ALETHEIA is that it incurs a low computational overhead.

THEOREM 3. *ALETHEIA runs in $O(N \log N + E + NKE)$, where N, K and E are the number of bidders, auctioned channels and edges in the conflict graph, respectively.*

5.3 Proofs

5.3.1 Proof of Theorem 1

One key observation is that, during the allocation process of ALETHEIA, the allocated channels always satisfy the conflict constraints. Therefore, it is sufficient for us to show for a bidder i with $b_i > p_i$, there are sufficient available channels for allocation to i. In the following, we prove this property by induction on N, the number of bidders.

Base Case. $N = 1$. In this case, the root node is the first node, and its request is less than (or equal to) K, can be allocated and thus the theorem holds.

Induction Step. Now suppose the theorem holds when $N < m$, we show that it still holds when $N = m$.

We prove this by contradiction. Since the bidders are considered in a sequential manner, according to the induction hypothesis, it is straightforward to show that for any bidder $i < m$ such that $b_i > p_i$, bidder i can get enough channels. So we only need to consider bidder m and $b_m > p_m$. By contradiction, we assume that bidder m with $b_m > p_m$ cannot be allocated. According to the allocation algorithm, bidder m's request to be rejected only happens when those channels have been allocated to its neighbors, i.e.,

$$|\cup_{j \in \mathbb{N}(m)} \mathcal{A}_j| + d_m > K. \tag{1}$$

Consider the pricing procedure for bidder m, if we first reserve d_m channels for bidder m, the above inequality indicates that there must exist a bidder $j \in \mathbb{N}(m)$, who originally wins, but cannot be allocated any more. Now two cases need to be considered. **Case I:** $t_m \leq t_j$. Since bidder j is m's losing neighbor and thus the price for m is no less than $d_m \cdot t_j$, *i.e.*, $p_m \geq d_m \cdot t_j$. This contradicts with $b_m > p_m$. **Case II:** $t_m > t_j$. Consider the computing of the price for bidder j. After reserving d_j channels for bidder j, bidder m would lose according to (1) and thus the price for j is no less than $d_j \cdot t_m$. By the inductive hypothesis, we know that bidder j is a winner, i.e., $b_j > p_j$. Combining the above results, we obtain $b_j > d_j \cdot t_m$ which contradicts with $t_m > t_j$. Therefore, Theorem 1 holds for $N = m$. The proof is accomplished.

5.3.2 Proof of Theorem 2

We first show that a bidder cannot increase its utility by using multiple identifiers.

LEMMA 1. *A bidder cannot increase its utility by submitting false-name bids, provided that all other bidders and their bids are same.*

PROOF. Suppose a bidder i uses two names i_1 and i_2 to participate the auction, and obtains x and y channels under the names i_1 and i_2, respectively. We show bidder i using a single identifier can obtain a total of $z = x + y$ channels with a utility that is no less than the sum of that obtained by i_1 and i_2.

[Case I: When names i_1 and i_2 are used.] We prove that their critical bidders have the same per-channel bid price, *i.e.*, $t_{c(i_1)} = t_{c(i_2)}$. The fundamental basis is that they both win and have the same conflicting neighbors except themselves. The claim holds when $c(i_1) = c(i_2)$. Otherwise, w.l.o.g., we assume $c(i_1) < c(i_2)$. For bidder i_1, by the procedure of finding critical bidder, we know i_1's allocation would cause $c(i_1)$ to lose, *i.e.,* $| \cup_{j \in \mathbb{N}(i_1), j \le c(i_1)} \mathcal{A}_j | + x > K$. Since $c(i_2) \in \mathbb{N}(i_1)$ and $c(i_1) < c(i_2)$, we obtain $| \cup_{j \in \mathbb{N}(i_1), j < c(i_2)} \mathcal{A}_j | + x > K$ which means i_1's allocation would cause $c(i_2)$ to lose. Therefore, we have $t_{c(i_1)} = max(t_{c(i_1)}, t_{c(i_2)})$.

Now consider computing price for bidder i_2, reserving y channels would cause $c(i_2)$ loses, but i_1 is i_2's winning neighbor since both i_1 and i_2 win. By the above deduction, i_1's allocation would cause $c(i_1)$ to lose and thus we have $t_{c(i_2)} = max(t_{c(i_1)}, t_{c(i_2)})$.

[Case II: When a single name i is used.] Since the requested channels in both cases are the same (*i.e.,* $z = x + y$) and i has the same conflict conditions as i_1 and i_2, i.e., satisfying i's request would cause $c(i_1)$ and $c(i_2)$ to lose , leading to $t_{c(i)} = max(t_{c(i_1)}, t_{c(i_2)})$.

Finally, we achieve that the bidder has the same utility in both cases. Similarly, when a bidder uses multiple names, the same deduction still holds. □

Then we prove that a bidder cannot increase its utility by submitting a cheating bid when using a single name.

LEMMA 2. *A bidder, using a single name, cannot increase its utility by submitting a cheating bid. (provided that all other bidders and bids are same)*

PROOF. Assume that bidder i's true valuation is v_i, which is different from its bid b_i, *i.e.,* $b_i \neq v_i$. We consider the following cases.

- Case I: Bidder i loses by bidding b_i and wins by bidding v_i. We have $u_i(v_i) > 0 = u_i(b_i)$.
- Case II: Bidder i loses by bidding both b_i and v_i, we have $u_i(v_i) = u_i(b_i) = 0$.
- Case III: Bidder i wins by bidding b_i and loses by bidding v_i. This case only happens when $t_{c(i)} > t_i$, and thus we have $u_i(v_i) = 0 > u_i(b_i)$.
- Case IV: Bidder i wins by bidding both b_i and v_i. Recall that the critical bidder of bidder i only depends on its demands and other bidders' demands. Since these parameters do no change when its bid changes and thus the critical bidders are the same for both cases. Therefore, we get $u_i(v_i) = u_i(b_i)$.

In summary, we have $u_i(v_i) \ge u_i(b_i)$ in all cases, which completes the proof. □

Finally, the false-name-proofness of ALETHEIA follows by Lemma 1 and Lemma 2 immediately.

5.3.3 Proof of Theorem 3

In the first step (§5.1.1), ALETHEIA uses a BFS procedure to construct the bidder list which takes $O(N + E)$ and also sorts each bidder's neighbor nodes with a complexity of $O(N \log N)$ in the worst case. Therefore this step takes $O(N + E + N \log N)$ time. In the second step (§5.1.2), for each bidder, ALETHEIA-Prices takes $2KE$ (each edge is updated at most 2 times) to update the available channel information of all bidders' neighboring bidders. Therefore it takes $O(NKE)$ time to compute the prices for all bidders. In the last step (§5.1.3), ALETHEIA-Allocation uses $O(N)$ to allocate channels to bidders and uses $2KE$ to update the availability of each bidder's channel, and hence its complexity is $O(N + KE)$. By summing up the cost of all three steps, the theorem holds directly.

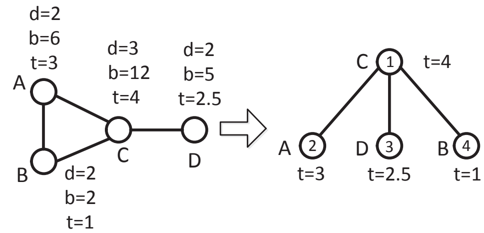

Figure 7: An illustration of demand reduction cheating, where bidder C can improve its utility when cheating by requesting 2 channels.

6. ALETHEIA EXTENSION

To make ALETHEIA more flexible and practical, we extend the basic ALETHEIA design, which only supports all-or-none requests, by enabling range-based requests. We name the enhanced design as ALETHEIA-RG, in which, bidders do not need to be single-minded. A bidder i can request d_i channels, and accepts any number of channels from 0 to d_i.

6.1 Valuation for Partial Satisfaction

Since requests of bidders are allowed to be partially satisfied, we need to rethink the corresponding valuation function. We make the following assumptions for the valuation function based on the fact that bidders are *rational*.

First, for each bidder i who requests d_i channels with true valuation $v(d_i)$, we assume that the valuation function satisfies *free disposal* [4]: for any $d_i' \ge d_i$, we have $v(d_i) = v(d_i')$. Free disposal indicates that when i's request is completely satisfied, allocating more channels will not improve its utility.

Second, we assume that the valuation function satisfies *super-additive* [4]: for any d' and d'', such that $d' + d'' \le d_i$ we have

$$v_i(d' + d'') \ge v(d') + v(d'). \quad (2)$$

Super-addition implies that a bidder would have higher utility when its request is closer to be fulfilled completely. Note that the above assumptions include the single-mind request as a special case.

6.2 New Cheating Pattern

When the requests can be partially satisfied, another cheating pattern, named *demand-reduction cheating* [6] arises, which may break the false-name-proofness of ALETHEIA. To see this, we introduce Figure 7 as an example. In this example, there are 4 bidders (A, B, C, and D) attend the auction to compete $K = 4$ channels. Their bids and demands, together with the conflict graph, are shown at the left side of Figure 7. According to ALETHEIA, bidder C wins and other bidders lose. Bidder C's critical bidder is A, and thus its utility is computed as $12 - 3 * 3 = 3$. Now we consider bidder C cheats on demand, by submitting $b_i = 8$ to request $d_i = 2$ channels, while other bidders stay unchanged. Now by the ALETHEIA, we know only bidder B loses and other bidders win, and C's critical bidder becomes bidder B. In this case, the utility of C is $8 - 1 * 2 = 6$. That is to say, bidder C can improve its own utility by reducing the demand.

Algorithm 3: Procedure $PreventDR()$

for $i = 1$ *to* N **do**
  if i *is a winner* **then**
    for *each* $< b_{i'}, d_{i'} > \in \mathcal{L}$ **do**
      $\mathcal{N} = \mathcal{N} \backslash \{i\} \cup \{i'\}$;
      $p_{i'} =$ ALETHEIA-Prices(i');
      $u_{i'} = v_{i'} - p_{i'}$;
      if $u_{i'} > u_i$ **then**
        $u_i = u_{i'}$;
        Re-allocate $d_{i'}$ channels to i.

6.3 Design of ALETHEIA-RG

To prevent demand reduction cheating, ALETHEIA-RG introduces a new component named PreventDR (Algorithm 3) besides the basic design ALETHEIA. In particular, after finding out the winners by running ALETHEIA, ALETHEIA-RG runs PreventDR for each winner to find the possible maximum utility when the winner reduces its channel demand to ensure the false-name-proofness.

PreventDR traverses all possible demand reduction cheatings, denoted as $d'_i \in \{1, 2, \ldots, d_i\}$, in order to find the maximal utility for each winner i in ALETHEIA. We denote $\mathcal{L} = \{\langle v_i(d_i), d_i \rangle, \langle v_i(d_i - 1), d_i - 1 \rangle, \ldots, \langle v_i(1), 1 \rangle\}$ as the cheating set. To ease the description, we use i' to denote the bidder i with cheating demand. For each possible cheating $\langle b_{i'}, d_{i'} \rangle \in \mathcal{L}$, we re-run ALETHEIA-prices by substituting i with i' to compute its new utility $u_{i'}$. If the newly obtained utility is greater than the obtained utility that i bids b_i for d_i channels, then Algorithm 3 withdraws previous allocation of i and reallocate $d_{i'}$ to i. Note that, the remaining channels $(d_i - d_{i'})$ will not be allocated to other bidders.

6.4 Properties with Proofs

THEOREM 4. *ALETHEIA-RG satisfies allocation feasibility.*

PROOF. Since ALETHEIA-RG is based on the results of ALETHEIA and it only decreases the number of allocated channels of each winner in ALETHEIA. Thus, ALETHEIA-RG does not violate the allocation feasibility property. □

THEOREM 5. *ALETHEIA-RG is false-name-proof.*

PROOF. The proof is very similar to the one for Theorem 2, and also consists of two parts. We thus only describe the differences between them.

For the first part (similar to Lemma 1), consider that bidder i uses two false names i_1 and i_2. The only differences coming from ALETHEIA-RG are that i_1 and i_2 may choose demand reduction cheating. Assume that the utility of bidder i_1 is maximized when bidding $b_{i'_1}$ for $d_{i'_1}$, and thus its utility is $u_{i'_1} = v_i(d_{i'_1}) - d_{i'_1} \cdot t_{c(i'_1)}$. Similarly, we get the utility of i_2 as $u_{i'_2} = v_i(d_{i'_2}) - d_{i'_2} * t_{c(i'_2)}$. Now we consider the same auction except that bidder i participates under a single name. If i chooses to lie on its demand, we have $d_{i'} \geq d_{i'_1}$ and $d_{i'} \geq d_{i'_2}$. Larger demands mean that we need to reserve more channels when computing the charging price, and thus more of its neighbors cannot be allocated, leading to $t_{c(i')} \geq t_{c(i'_1)}$ and $t_{c(i')} \geq t_{c(i'_2)}$. Moreover, by Eq. (2), we get $v_i(d_{i'_1} + d_{i'_2}) \geq v_i(d_{i'_1}) + v_i(d_{i'_2})$. Combining these above equations, we obtain

$$u_{i'} = v(d_{i'_1} + d_{i'_2}) - (d_{i'_1} + d_{i'_2}) \cdot t_{c(i')} \geq u_{i'_1} + u_{i'_2}. \quad (3)$$

In other words, if demand reduction cheating is profitable for bidder i, our claim holds by equation (3). If demand reduction cheating is not profitable for bidder i, we have $u_i \geq u_{i'} \geq u_{i'_1} + u_{i'_2}$, again our claim holds.

For the second part (similar to Lemma 2), the differences also come from the case that bidder i chooses to make demand reduction cheating. Such cheating only happens when i is a winner, and its payment is irrelevant to b_i and v_i $(b_i \neq v_i)$. Thus it achieves the same utility no matter whether it cheats on the demand or not. Therefore, we conclude that a bidder using a single name cannot increase its utility by cheating on bids. □

THEOREM 6. *ALETHEIA-RG runs in polynomial complexity of $O(N^2 K \log N + N^2 K^2 E)$.*

PROOF. For each winner i, we need to run ALETHEIA-prices at most d_i times and thus the prevent demand reduction lie procedure takes $N \cdot d_i \cdot O(N \log N + NKE)$ for all winners. Therefore, the ALETHEIA-RG runs in complexity less than $O(N^2 K \log N + N^2 K^2 E)$. □

7. EVALUATION

7.1 Simulation Setup

We randomly deploy bidders in a 100×100 square area, and set the interference range as 10. If the distance between any two bidders is less than 10, they interfere with each other when accessing the same channel at the same time. The per-channel bids of bidders are randomly distributed in the range of [0, 1]. We run the simulation 100 times and average the results to get rid of the randomness and observe the general patterns.

7.2 Performance

We use the *revenue* and *spectrum utilization* as our performance metrics, which are defined as the total payments and numbers of allocated channels of all winners, respectively. We compare ALETHEIA with the simple strawman approach proposed in §4.2. For ease description, we name it as SIMPLE.

We compare them in two scenarios. First, we set the number of bidders as $N = 300$ and vary the number of auctioned channels. Each bidder requests either 1 or 2 channels. The results are plotted in Figure 8. ALETHEIA and SIMPLE achieve similar peak revenue when 4 and 8 channels are used, respectively, as illustrated in Figure 8(a). Figure Figure 8(b) shows that ALETHEIA significantly improves spectrum utilizations compared to SIMPLE under any number of auctioned channels. Figure 8(c) plots revenue as a function of spectrum utilization. Again, ALETHEIA performs much better. Second, we compare them by varying the number of bidders, setting the default number of auctioned channels as $K = 4$, and the request of each bidder is randomly draw from set $\{1, 2, 3, 4\}$. The results are shown in Figure 10. In this case, ALETHEIA significantly improves revenue and spectrum utilization compared to SIMPLE by up to 300%. This is because SIMPLE sacrifices spectrum reuse to maintain false-name-proofness. When the number of bidders grows, the number of allocated bidders in each partitioned

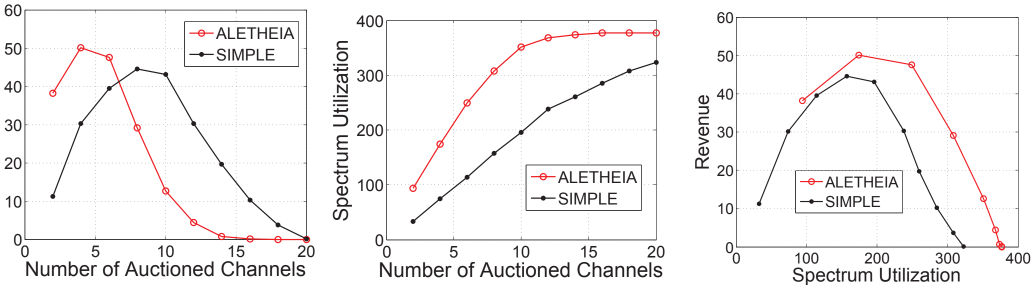

(a) Revenue (b) Spectrum Utilization (c) Revenue vs. Spectrum Utilization

Figure 8: Compare ALETHEIA with SIMPLE by varying the number of auctioned channels.

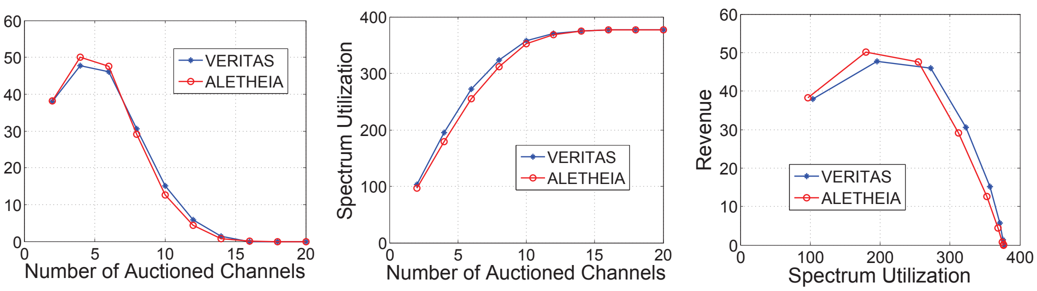

(a) Revenue (b) Spectrum Utilization (c) Revenue vs. Spectrum Utilization

Figure 9: Compare ALETHEIA with VERITAS by varying the number of auctioned channels.

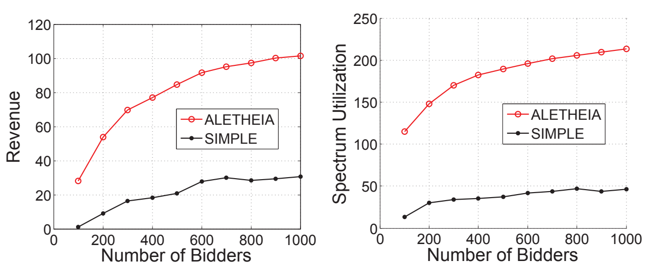

(a) Revenue (b) Spectrum Utilization

Figure 10: Compare ALETHEIA with SIMPLE by varying the number of bidders.

box in SIMPLE is nearly constant. In contrast, ALETHEIA can accommodate more bidders, i.e., achieving much higher spectrum reuse.

7.3 Cost of False-name-proofness

We also compare ALETHEIA with VERITAS to evaluate the cost of false-name-proofness. Specifically, we try to answer the question: how much performance would ALETHEIA lose if no bidders submit false-name bids?

We compare them by vary the number of auctioned channels from 2 to 20 and set the number bidders as 300. Each bidder's request is either 1 or 2. The results are shown in Figure 9. We can see that ALETHEIA has very similar performance, in terms of both revenue and utilization, as VERITAS. In order to guarantee false-name-proofness, ALETHEIA sacrifices part of bidders who have no critical neighbors in VERITAS (as shown in Section 3.1) and thus it performs slightly worse than VERITAS. On other hands, when the competition between bidders is high (i.e., when the number of auctioned channels is small), ALETHEIA performs slightly better than VERITAS in terms of revenue. This is because ALETHEIA charges by the highest losing neighbor's per-channel bid while VERITAS charges by that of the first losing neighbor.

8. RELATED WORK

Spectrum allocation mechanisms have been extensively studied in recent years. A number of auction designs have been proposed to improve spectrum utilization and allocation efficiency. VERITAS [24] is one of the pioneer auction designs with strategy-proofness that exploits the spectrum reusability in radio spectrum. Later, the work is extended to consider double spectrum auctions [25]. Jia *et al.* [9] and Al-Ayyoub *et al.* [3] design spectrum auctions to maximize the expected revenue by assuming that the bids of the secondary users follow certain distributions. SMALL [20] is designed for the scenario where the primary user sets a reserved price for each channel. TRUMP [19] allocates spectrum access rights on the basis of QoS demands. The aforementioned work achieve strategy-proofness under the assumption that each bidder is associated with a unique name. As a result, none of them can resist false-name bid cheating where bidders can improve their utilities by submitting multiple false-name bids simultaneously. Another group of work (e.g., [21, 26]) study the collusion issue among bidders in spectrum auctions. Compared with collusion, while a bidder has to seek out and persuade other bidders to join in collusion, a false-name bid is easier to make since it can be done alone.

With the popularity and success of Internet auctions, false-name bid cheating has attracted more and more research interests. The effects of false-name bids on combinatorial auctions are analyzed in [23]. Following that, a series of false-name-proof mechanisms have been developed for a variety of scenarios. For example, Yokoo *et al.* [12] designed a leveled division set based mechanism for multi-item single-unit auctions and prove its false-name-proofness. In [11,17], false-name-proof mechanisms have been proposed to address the multi-unit single-item settings. Unfortunately, all these designs are for traditional goods (e.g., books or paintings). They do not consider the spectrum reusability and cannot deal with the complex inference constraints in spectrum auctions. As a result, when being applied to spectrum auctions, these designs either lose false-name-proofness or create excess interference, as discussed in Section 4.

Different from the previous studies, ALETHEIA not only provides strategy-proofness but also resists false-name bid cheating. To the best of our knowledge, ALETHEIA is the first spectrum auction design that can achieve false-name-proofness.

9. CONCLUSIONS

In this paper, we studied the new type of cheating named false-name bids, in large-scale spectrum auctions. We demonstrated that false-name bid cheating is easy to form in existing strategy-proof spectrum auctions and can severely impair the auction revenue. We further devised ALETHEIA, the first false-name-proof design for spectrum auction that nullifies the possibility of increasing profits by submitting false-name bids. It achieves high spectrum redistribution efficiency and low computational overhead, and is very flexible to support diverse request formats of bidders.

10. REFERENCES

[1] http://www.atheros.com.

[2] I. Akyildiz, W. Y. Lee, M. Vuran, and S. Mohanty. Next generation/dynamic spectrum access/cognitive radio wireless networks: a survey. *Computer Networks*, 50(13):2127–2159, 2006.

[3] M. Al-Ayyoub and H. Gupta. Truthful spectrum auctions with approximate revenue. In *Proc. of INFOCOM*. IEEE, 2011.

[4] Y. Bartal, R. Gonen, and N. Nisan. Incentive compatible multi unit combinatorial auctions. In *Proc. of TARK*, 2003.

[5] T. Clancy and N. Goergen. Security in cognitive radio networks: Threats and mitigation. In *Proc. of CrownCom*. IEEE, 2008.

[6] R. Engelbrecht-Wiggans, J. List, and D. Reiley. Demand reduction in multi-unit auctions with varying numbers of bidders: Theory and field experiments. *Inernational Economic Review*, 47(1):203–231, 2006.

[7] K. Ezirim, E. Troja, and S. Sengupta. Sustenance against rl-based sybil attacks in cognitive radio networks using dynamic reputation systems. In *Proc. of MILCOM*. IEEE, 2013.

[8] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu. Impact of interference on multi-hop wireless network performance. In *Proc. of MOBICOM*. ACM, 2003.

[9] J. Jia, Q. Zhang, Q. Zhang, and M. Liu. Revenue generation for truthful spectrum auction in dynamic spectrum access. In *Proc. of MOBIHOC*. ACM, 2009.

[10] P. Klemperer. *Auctions, theory and practice.* Princeton University Press, USA, 2004.

[11] S. M. M. Yokoo, Y. Sakurai. Robust multi-unit auction protocol against false-name bids. In *Proc. of IJCAI*. IEEE, 2001.

[12] S. M. M. Yokoo, Y. Sakurai. Robust combinatorial auction protocol against false-name bids. *Artificial Intelligence*, 130(1):167–181, 2002.

[13] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory.* Oxford Press, USA, 1995.

[14] M. McHenry, P. Tenhula, D. McCloskey, D. Roberson, and C. Hood. Chicago spectrum occupancy measurements & analysis and a long-term studies proposal. In *Proc. of TAPAS*, 2006.

[15] M. J. Osborne and A. Rubenstein. *A Course in Game Theory.* MIT Press, USA, 1994.

[16] Y. Tan, S. Sengupta, and K. Subbalakshmi. Spectrum stealing via sybil attacks in dsa networks: implementation and defense. In *Proc. of ICC*. IEEE, 2011.

[17] K. Terada and M. Yokoo. False-name-proof multi-unit auction protocol utilizing greedy allocation based on approximate evaluation values. In *Proc. of AAMAS*, 2003.

[18] T. Todo, A. Iwasaki, M. Yokoo, and Y. Sakurai. Characterizing false-name-proof allocation rules in combinatorial auctions. In *Proc. of AAMAS*, 2009.

[19] Q. Wang, B. Ye, S. Lu, and S. Guo. A truthful qos-aware spectrum auction with spatial reuse for large-scale networks. *TPDS*, 46(1):2499–2508, 2014.

[20] F. Wu and N. Vaidya. Small: A strategy-proof mechanism for radio spectrum allocation. In *Proc. of INFOCOM*. IEEE, 2011.

[21] Y. Wu, B. Wang, K. Liu, and T. Clancy. A scalable collusion-resistant multi-winner cognitive spectrum auction game. *Transactions on Communications*, 57(12):3805 – 3816, 2009.

[22] M. Yokoo, Y. Sakurai, and S. Matsubara. The effect of false-name declarations in mechanism design: Towards collective decision making on the internet. In *Proc. of ICDCS*. IEEE, 2000.

[23] M. Yokoo, Y. Sakurai, and S. Matsubara. The effect of false-name bids in combinatorial auctions: New fraud in internet auctions. *Games and Economic Behavior*, 46(1):174–188, 2004.

[24] X. Zhou, S. Gandhi, S. Suri, and H. Zheng. ebay in the sky: Strategy-proof wireless spectrum auctions. In *Proc. of MOBICOM*. ACM, 2008.

[25] X. Zhou and H. Zheng. Trust: A general framework for truthful double spectrum auctions. In *Proc. of INFOCOM*. IEEE, 2009.

[26] X. Zhou and H. Zheng. Breaking bidder collusion in large-scale spectrum auctions. In *Proc. of MOBIHOC*. ACM, 2010.