

eBay in the Sky: Strategy-Proof Wireless Spectrum Auctions

Xia Zhou, Sorabh Gandhi, Subhash Suri and Haitao Zheng
Department of Computer Science
University of California, Santa Barbara, CA, USA
{xiazhou, sorabh, suri, htzheng}@cs.ucsb.edu

ABSTRACT

Market-driven dynamic spectrum auctions can drastically improve the spectrum availability for wireless networks struggling to obtain additional spectrum. However, they face significant challenges due to the fear of market manipulation. A truthful or strategy-proof spectrum auction eliminates the fear by enforcing players to bid their true valuations of the spectrum. Hence bidders can avoid the expensive overhead of strategizing over others and the auctioneer can maximize its revenue by assigning spectrum to bidders who value it the most. Conventional truthful designs, however, either fail or become computationally intractable when applied to spectrum auctions. In this paper, we propose VERITAS, a truthful and computationally-efficient spectrum auction to support an eBay-like dynamic spectrum market. VERITAS makes an important contribution of maintaining truthfulness while maximizing spectrum utilization. We show analytically that VERITAS is truthful, efficient, and has a polynomial complexity of $O(n^3k)$ when n bidders compete for k spectrum bands. Simulation results show that VERITAS outperforms the extensions of conventional truthful designs by up to 200% in spectrum utilization. Finally, VERITAS supports diverse bidding formats and enables the auctioneer to reconfigure allocations for multiple market objectives.

Categories and Subject Descriptors

C.2.1 [Computer Systems Organization]: Computer-Communication Networks

General Terms

Algorithm, Design, Economics

Keywords

Spectrum Auctions, Mechanism Design

1. INTRODUCTION

An increasing number of users, homes and enterprises rely on wireless technology for their daily activities. However,

the growth of wireless networks has been hampered by the inefficient distribution of radio spectrum. Historical static allocations have led to an artificial shortage of spectrum: new wireless applications starve for spectrum, while large chunks of it remain idle most of the time under their current owners. This misallocation has prompted a wide interest in an open, market-based approach for redistributing the spectrum where new users can gain access to the spectrum they desperately need and existing owners can gain financial incentives to “lease” their idle spectrum.

Auctions are among the best-known market-based allocation mechanisms due to their perceived fairness and allocation efficiency – everyone has an equal opportunity and the goods are sold to bidders who value them the most. Indeed, FCC (Federal Communications Commission) and its counterparts across the world have auctioned unused spectrum for billions of dollars in the past decade. However, a FCC-style spectrum auction targets long-term national/regional leases, requiring huge up-front investments. It often takes months or years to conclude, involves only a few large corporate players, and entails significant manual negotiations.

In this paper, we introduce a very different auction format to support the open market-based spectrum redistributions. We consider a dynamic spectrum auction system akin to the eBay marketplace that serves and scales to many small players without manual mediations. In this marketplace, wireless nodes request spectrum in their *local* neighborhoods in *short-terms*. These small players request spectrum based on present demands and pay for what they really need without burdensome up-front investments in the FCC-style auctions.

One critical requirement to initiate the proposed marketplace is to ensure that auctions are quickly conducted to enable on-demand short-term spectrum redistribution. Bidders request spectrum on the fly and the auctioneer processes them immediately. From an economic perspective, the fundamental obstacle is the significant overhead taken by both the auctioneer and bidders to avoid market manipulation. Bidders must strategize over others on how to bid, and the auctioneer applies Bayesian settings to increase its revenue. This overhead and the fact that it might not be the best strategy can easily discourage bidders from participations. Consequently, an auction with too few bidders will be both unprofitable for the auctioneer and potentially inefficient.

In response to this challenge, we develop a truthful (or strategy-proof) spectrum auction. A truthful auction guarantees that if a bidder bids the true valuation of the resource, its utility will not be less than that when it lies.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiCom '08, September 14–19, 2008, San Francisco, California, USA.
Copyright 2008 ACM 978-1-60558-096-8/08/09 ...\$5.00.

Hence, the dominating strategy for a bidder is to bid its true valuation. To bidders, a truthful auction eliminates the expensive overhead of strategizing about other bidders and prevents market manipulation. Thus it can attract a wide range of network nodes/establishments to engage in the marketplace. To the auctioneer, by encouraging bidders to reveal their true valuations, a truthful auction can help the auctioneer increase its revenue by assigning spectrum to the bidders who value it the most. For the same reason, many classical auction systems are made truthful, including the sealed-bid secondary-price [15], k-position [10, 11] and VCG auctions [2, 5].

While prior works have enforced truthfulness in conventional auctions, we show that existing truthful designs either fail or become computationally prohibitive when applied to spectrum auctions. The fundamental reason is that unlike goods (*e.g.* paintings, bonds, electricity) in conventional auctions, spectrum is *reusable* among bidders subjecting to the spatial interference constraints. Because interference is only a local effect, bidders in close proximity cannot use the same spectrum frequency simultaneously but well-separated bidders can. These heterogeneous interdependencies among bidders make secondary-price and k-position auctions no longer truthful. Furthermore, these constraints make the problem of finding the optimal spectrum allocation NP-complete [8], and hence a real-time spectrum auction with many bidders must resort to greedy allocations that are computationally efficient. Unfortunately, it has been shown that the VCG auction loses its truthfulness under greedy allocations [9].

In this paper, we propose VERITAS, a truthful dynamic spectrum auction framework that only requires polynomial complexity. VERITAS is a sealed-bid truthful auction: bidders submit their bids privately to the auctioneer, who grants spectrum channels to the selected bidders but charges them with prices equal or less than their actual bids. VERITAS provides the following key advantages:

- VERITAS achieves truthfulness with computationally-efficient spectrum allocation and pricing mechanisms, making it feasible for the online short-term auction.
- VERITAS provides the auctioneer with the capability and flexibility of maximizing its customized objective. The auctioneer can configure the order of allocation and the amount of spectrum offered to maximize the revenue or the social welfare.
- VERITAS provides the bidders the flexibility of diverse demand formats. A bidder can request spectrum by the exactly number of channels it would like to obtain, or by a range defined by the minimal and maximal number of channels.

Our analytical and experimental results reveal the following findings:

- VERITAS provides truthful and computationally-efficient spectrum auctions by sequentially allocating spectrum to bidders following a bid-dependent ranking, and by charging each winner with the minimum it needs to pay to win the auction.
- VERITAS has a computational complexity of $O(n \log n + nk|E|)$, where n is the number of bidders, k is the number of channels auctioned, and $|E|$ is the number

of edges in the conflict graph used to represent the interference condition among bidders.

- The auctioneer can customize the ranking metric to achieve desired market outcomes, such as maximizing the revenue. VERITAS's spectrum allocation algorithm performs similarly to the best-known greedy allocation algorithms that do not consider truthfulness.
- In untruthful auctions, the auction revenue increases with the number of winning bidders. In contrary, the revenue of truthful auctions (and VERITAS) increases with the number of winning bidders initially but decreases when the winning bidders exceed a threshold. VERITAS introduces a screening mechanism that optimizes the number of channels auctioned to maximize the revenue.

2. PRELIMINARIES AND PROBLEM DEFINITION

We start by introducing a set of notations used to define a truthful spectrum auction. We consider a collusion-free spectrum auction setting, where one auctioneer auctions k channels to n bidders located in a geographic region. We assume that the channels have uniform characteristics and values, so that bidders request spectrum by submitting the number of channels they demand and the per-channel prices they would like to pay. To make the problem tractable, we represent the conflict condition among bidders by a conflict graph – two bidders either interfere with each other and cannot use the same channels, or can reuse the same channels simultaneously.

Channel request (d_i) – It represents the number of channels requested by bidder i . In *strict* requests, a bidder accepts to receive either d_i channels or 0 channel; in *range* requests, a bidder accepts any x channels if $0 \leq x \leq d_i$.

Per-channel bid (b_i) – It represents the per-channel bid submitted by bidder i . Let $B = \{b_1, b_2, \dots, b_n\}$ represent the set of bids submitted by all the bidders.

Per-channel true value (v_i) – Each bidder i has a per-channel valuation v_i which describes the true price i is willing to pay for each channel. In most cases, this valuation is private and is known only to the bidder itself.

Clearing price (p_i) – Given the bid set B , the auctioneer will allocate channels to bidders and charge price p_i , referred to as the *clearing price* for each winner i . This price might be different across bidders but must not exceed the bid submitted by i times the number of channels assigned to i .

Bidder utility (u_i) – The *utility* of bidder i is the residual worth of the channels. That is, $u_i = v_i \cdot d_i^a - p_i$ if i obtains d_i^a channels, and 0 if it obtains none.

DEFINITION 1. *A truthful auction is one in which no bidder i can obtain higher utility u_i by setting $b_i \neq v_i$.*

DEFINITION 2. *An efficient and a truthful spectrum auction is one which is truthful and maximizes the efficiency of spectrum usage subject to the interference constraints.*

Given the above definitions, we now describe the two unique properties that set spectrum auctions fundamentally different from (and much more difficult than) conventional multi-unit auctions. First, spectrum can be spatially reused

concurrently – two conflicting bidders must not use the same channels simultaneously yet well-separated bidders can. While a conventional auction with n bidders and k channels can only have at most k winners, spectrum auction can have more than k winners. Let’s consider a simple example of $n = 3$ bidders competing for $k = 2$ channels, each requesting 1 channel. Figure 1(left) plots the conflict graph, where each vertex represents a bidder and two vertices share an edge if they conflict. A conventional auction will sell channels to at most 2 bidders, while an *efficient* spectrum auction can assign channels to all 3 bidders. Second, the conflict constraints among bidders are in general heterogenous, making the problem of optimizing spectrum allocation NP-complete [8] even when each bidder requests one channel.




Figure 1: An illustrative example on spectrum allocations. (Left) The conflict graph of a network with 3 bidders. (Right) The optimal spectrum allocation when there are 2 channels.

Next, we show that these unique properties of spectrum allocation bring significant challenges into truthful and efficient spectrum auction designs. Existing truthful designs in conventional auctions, when applied to spectrum auctions, either fail to be truthful, require exponential computational complexity, or significantly degrade spectrum utilization.

3. CHALLENGES OF TRUTHFUL AND EFFICIENT SPECTRUM AUCTION DESIGN

In this section, we illustrate the challenges in designing truthful and efficient spectrum auctions. We start from two truthful designs from conventional auctions and show that they become untruthful when applied to spectrum auctions. We then introduce a simple (naive) design of truthful spectrum auction which leads to significant loss of spectrum usage, but serves as the baseline in our paper. For the ease of understanding, we assume each bidder requests one channel. The same conclusion applies to the scenario where bidders request multiple channels. In each auction design, the pricing is applied after the spectrum allocation.

3.1 Secondary Pricing Spectrum Auctions

Consider the following auction algorithm. Sort the bids in descending order. Allocate one channel to each of the top l bidders and charge them the $(l + 1)$ th bidder’s bid. This secondary pricing auction was originally proposed and shown to be truthful in the seminal paper of Vickery [15], and later extended to cases where each bidder requests more than one items [10, 11]. The natural extension to spectrum auctions leads to the following allocation and pricing algorithms.

Allocation:

- Sort the bids in descending order and set each bidder’s available channel set as channel 1 to k .
- Allocate a channel m to the first bidder i in the sorted order using the lowest indexed channel in i ’s available channel set, remove i from the list, remove m from i ’s conflicting neighbors’ available channel sets.




Figure 2: This example shows why the one with secondary pricing is untruthful. The left part shows the auction results when all bidders truthfully bid. The right part shows that bidder c improves its utility by bidding higher than its true value.




Figure 3: This example shows why a VCG-style spectrum auction is untruthful. When bidder c raises its bid (right part), it increases its utility by getting a channel and paying less than its true value.

- Repeat 2 until all the bidders have been considered.

Pricing:

Charge winner i the highest bid of its unallocated conflicting neighbors. If there is no such neighbor, charge 0.

We show that the above auction is untruthful using a counter example. Figure 2 shows the conflict graph of 4 bidders (a, b, c, d) competing for 2 channels. When bidding truthfully, the utilities are 5, 3, 0, 1 respectively. However, when bidder c cheats by raising its bids to 3, it will obtain a channel and be charged with 0, increasing its utility to 1. Hence, by bidding untruthfully bidder c has improved its utility, which contradicts the definition of truthfulness.

3.2 VCG-Style Spectrum Auctions

We now consider applying VCG-style designs to spectrum auctions. VCG auctions are generally intractable because they require the optimal allocation of resources. Existing works have exploited greedy allocation schemes paired with the VCG-style pricings to design truthful auction algorithms [10, 11]. For spectrum auction with combinatorial interference constraints, the natural extension (with polynomial complexity) is to use a greedy allocation (of Section 3.1) with a VCG-style pricing mechanism. The price charged to bidder i is the bid of its first rejected neighbor who would have been allocated if i were absent from the auction.

Again, we use a counter example to show that this design is untruthful in the context of spectrum auctions. Figure 3 shows a network of 5 bidders. When bidder c changes its bid from 2 to 4, it increases its utility from 0 to 1.

3.3 A Simple Truthful Spectrum Auction

Next we show that when the corresponding conflict graph is a unit disk graph, existing truthful designs can be applied

to spectrum auctions by sacrificing spectrum utilization significantly. For simplicity, we assume that the bidders are located in a rectangular region. We divide the region into squares, each with the length of the maximal interference range of the bidders, and split k channels into 4 subsets with $k/4$ channels each. In each square box, we can apply the original secondary pricing mechanism assuming all the bidders in this box conflict with each other. That is, we allocate $k/4$ channels to the top $k/4$ bidders in each box and charge them the bid of the $(k/4 + 1)$ th bidder in the same box. It is straightforward to show that this auction design is truthful, and hence we omit the proof.

While being truthful, this auction design suffers from significant degradation in spectrum utilization. The static partition of spectrum among the boxes leads to at least a factor of 4 in spectrum degradation, let alone the degradation within each box where bidders are assumed to conflict with all others. This observation, in fact, motivates us to design a sophisticated spectrum auction that achieves truthfulness and utilizes spectrum efficiently.

4. VERITAS AUCTION DESIGN

Motivated by the observations from Section 3, we propose VERITAS, a *truthful* and *computationally-efficient* spectrum auction design that also *utilizes spectrum efficiently*. VERITAS consists of a greedy spectrum allocation algorithm to distribute channels among bidders and a pricing mechanism to charge winning bidders. By strategically designing the greedy allocation algorithm, VERITAS achieves similar spectrum utilization/efficiency as the well-known spectrum allocation algorithms in polynomial time. By designing a pricing mechanism to charge each winner with the minimum it needs to pay to win the auction, VERITAS enforces auction truthfulness despite the complex heterogeneous interference constraints.

We design VERITAS to support diverse forms of spectrum requests. In this section, we introduce the main algorithm of VERITAS with *strict requests*, and the proofs of its truthfulness and computational complexity. In strict requests, a bidder i requests spectrum by d_i channels and only accepts allocations of either 0 or d_i channels. We show in Section 5 that VERITAS can be easily extended to three other bidding formats, namely (i) *range requests* where bidder i requests spectrum by d_i channels but accepts to receive any number of channels between 0 and d_i , (ii) *contiguous strict requests* where the channels in strict requests assigned to i must be contiguously aligned, and (iii) *contiguous range requests* where the channels in range requests assigned to i must be contiguously aligned.

4.1 VERITAS Main Algorithm

We start from the main VERITAS algorithm designed for strict requests. We represent each bidder i 's bid as $[d_i, b_i]$, where d_i is the number of channels requested by i and b_i is the per-channel bid from i . We describe VERITAS by its spectrum allocation and pricing algorithms. We first assume that the bid set B is sorted in descending order of b_i . In Section 4.3 we show that VERITAS can use flexible sorting functions $f(b_i)$ and only require $f(b_i)$ to be an increasing function of the bid and not affected by other bidders.

For easy illustration, we first introduce a few notations.

- $D = \{d_1, d_2, \dots, d_n\}$ represents the number of channels demanded by the bidders.

- $G = (V, E)$ represents the conflict graph where V is the collection of the bidders and E is the collection of edges where two bidders share an edge if they conflict.
- $N(i)$ represents the set of i 's conflicting neighbors who cannot share the same channel with i simultaneously, *i.e.* the set of bidders sharing edges with i in G .
- $\text{Distinct}(N(i))$ represents the distinct set of channels that have been assigned to all the members of $N(i)$.
- $\text{TOP}(B)$ represents the 1st bidder in the bidding set B sorted in a manner of descending b_i .
- $a_{ij} = 1$ if channel j is assigned to bidder i , else 0.

VERITAS–Allocation Based on the sorted bid set B , the algorithm (see Algorithm 1) allocates bidders sequentially from the highest one to the lowest one. For each bidder i , the algorithm first checks whether there are enough channels to satisfy i , *i.e.* $\text{Distinct}(N(i)) + d_i \leq k$. If so, the function $\text{Assign}(i, d_i)$ assigns to i d_i channels with the lowest available indices that are not in $\text{Distinct}(N(i))$.

Algorithm 1 VERITAS-Alloc(B, D, G)

```

1:  $B'$  = sorted  $B$ 
2: while  $B' \neq \emptyset$  do
3:    $i$  =  $\text{TOP}(B')$ ;
4:   if  $\text{Distinct}(N(i)) + d_i \leq k$  then
5:      $\text{Assign}(i, d_i)$ 
6:   end if
7:    $B' = B' \setminus \{b_i\}$ 
8: end while

```

VERITAS–Pricing VERITAS charges each winner i with the bid of its *critical neighbor* multiplied by the number of channels allocated to i .

DEFINITION 3. Given $\{B \setminus b_i\}$, a *critical neighbor* $\mathbb{C}(i)$ of bidder i is a bidder in $N(i)$ where if i bids lower than $\mathbb{C}(i)$, i will not be allocated, and if i bids higher than $\mathbb{C}(i)$, i will be allocated.

At the first sight, finding the critical neighbor seems computationally expensive. It requires inserting i immediately after each of its neighbors and running Algorithm 1 repeatedly to determine i 's allocation statuses. VERITAS overcomes this problem by introducing an intelligent pricing algorithm that identifies the critical neighbor for each bidder by running Algorithm 1 only once. The basic idea is that for each bidder i , first take i out of the sorted bid list B' and run VERITAS-Alloc. When assigning channels to i 's neighbors, remove the allocated channels from i 's available channel set. The first winning neighbor who makes the number of i 's available channels below its demand d_i is i 's critical neighbor. We describe the VERITAS-Pricing algorithm in Algorithm 2, where *avail_ch* is the set of currently available channels at i , *owned_ch* is the number of channels currently owned by i , and $\{c_1, c_2, \dots, c_k\}$ is the set of k channels.

Toy Example Consider the example in Figure 3 and take bidder c as an instance. When $k = 2$, c is denied when bidding truthfully, resulting in zero utility. When c raises its bid to 4, VERITAS charges c by its critical neighbor e 's bid which is 3, and c 's utility becomes $2 - 3 = -1 < 0$. Therefore, c cannot improve its utility by bidding untruthfully.

Algorithm 2 VERITAS-Pricing(B, D, G, i)

```
1:  $B' = \text{sorted } B$ 
2:  $\text{owned\_ch} = \sum_{j=1}^k a_{ij}$ 
3: if  $\text{owned\_ch} = 0$  then
4:    $p_i = 0$ 
5:   return
6: end if
7:  $p_i = 0$ 
8:  $\text{avail\_ch} = \{c_1, c_2, \dots, c_k\}$ 
9:  $B'' = B' \setminus \{b_i\}$ 
10: while ( $B'' \neq \emptyset$ ) AND ( $\text{owned\_ch} > 0$ ) do
11:    $q = \text{TOP}(B'')$ 
12:   if  $|\text{Distinct}(N(q))| + d_q \leq k$  then
13:      $\text{Assign}(q, d_q)$ 
14:     if  $q \in N(i)$  then
15:        $\text{avail\_ch} = \text{avail\_ch} \setminus \text{channels allocated to } q$ 
16:       if  $|\text{avail\_ch}| < \text{owned\_ch}$  then
17:          $p_i = b_q \cdot d_i$ 
18:          $\text{owned\_ch} = 0$ 
19:       end if
20:     end if
21:   end if
22:    $B'' = B'' \setminus \{b_q\}$ 
23: end while
```

Next, we prove that VERITAS is truthful. We show that the combination of the greedy allocation and critical neighbor based pricing is essential to achieve truthfulness.

4.2 VERITAS Truthfulness

We prove that VERITAS is truthful by showing that given B in one private auction round, no bidder can obtain higher utility by bidding other than its true value. Our proof consists of three steps: (1) We prove that VERITAS's allocation is *monotonic* – given B , if a bidder is allocated by bidding b , then it will also be allocated by bidding higher than b . (2) We show that for each bidder, there exists a *critical value* such that the bidder wins by bidding higher than this value and loses by bidding lower. The critical value is the bid of its critical neighbor. (3) We show that by charging winning bidders based on their respective critical values, no bidder can obtain higher utility by bidding other than its true value.

(1) Monotonic Allocation

LEMMA 1. *If any bidder i is allocated by bidding b_i^1 , it will also be allocated if it bids b_i^2 , where $b_i^2 > b_i^1$ (provided all the other bids and channel demands remain the same).*

PROOF. We prove this lemma by contradiction. Consider two sorted lists of bids, B_1 and B_2 in Figure 4. The bids of B_1 and B_2 are the same except for bidder i . In B_1 , bidder i bids b_i^1 , and in B_2 , b_i^2 . Define the position of bidder i in B_1 and B_2 by $\text{pos}(b_i^1)$ and $\text{pos}(b_i^2)$ respectively. Since $b_i^1 < b_i^2$, we have $\text{pos}(b_i^1) > \text{pos}(b_i^2)$. Moreover, the two sorted bid lists are exactly the same before $\text{pos}(b_i^2)$ because the demands and bids of all other bidders remain the same. Assume that bidder i is not allocated with B_2 , then the number of available channels at $\text{pos}(b_i^2)$ is less than d_i . But VERITAS-Alloc algorithm never deallocates channels, hence there must be strictly less than d_i channels for all positions after $\text{pos}(b_i^2)$ in B_1 , which implies that bidder i must not be allocated in B_1 . This is a contradiction. Hence it cannot be that bidder i is allocated in B_1 and not allocated in B_2 . \square




Figure 4: Two sorted bid lists where only i 's bid is different. Bids between positions X and T are the same for both list B_1 and B_2 , also bids between positions W and Z are the same for both lists.

Similarly, we can also show the following:

LEMMA 2. *If any bidder i is rejected by bidding b_i^2 , its bid will also be rejected if it bids b_i^1 , where $b_i^2 > b_i^1$ (provided all the other bids and channel demands remain the same).*

(2) Critical Neighbor/Value

The two lemmas above lead us to the following key lemma.

LEMMA 3. *For any bidder i , if i would be rejected by bidding some value, then there exists a unique position in the sorted bid list, such that if i 's bid is placed before that position i will win, and if i 's bid is placed after that position it will lose. Moreover, that position is occupied by one of i 's neighbors in $N(i)$.*

PROOF. The proof for the first part simply follows Lemma 1 and Lemma 2. The number of available channels for any bidder i only decreases when the allocation is conducted for each bidder along the sorted bid list. And this can only happen when those channels are allocated to i 's neighbors, hence that unique position must be occupied by a neighbor of bidder i . \square

We define the bid of the neighbor in this critical position as the *critical value* of bidder i . When i can be allocated at any position in the bid list, then its critical value is 0. Based on the description of the VERITAS-Pricing algorithm (Algorithm 2), we arrive at the following lemma:

LEMMA 4. *VERITAS-Pricing charges winning bidders by their critical values multiplied by the number of channels they obtained, and losing bidders by zeros.*

Now we establish the second key lemma, which shows that if any bidder i wins, then its clearing price is no more than its submitted bid times the number of its assigned channels.

LEMMA 5. *For each winner i in VERITAS, its clearing price is less than (or equal to) its submitted bid b_i multiplied by the number of requested channels d_i .*

PROOF. There are two cases. If i can always be allocated, VERITAS charges it with zero which is equal to its critical value of zero. Our claim holds. Otherwise, VERITAS charges i based on the bid of its critical neighbor. Since i does obtain channels with b_i , it implies (by Lemma 1) that any bid which puts bidder i before its current position in the bid order (any per-channel bid higher than b_i) will also be satisfied. So the per-channel clearing price must either be derived from a bidder with bid $\leq b_i$ or 0, and hence i 's clearing price cannot be greater than $b_i \cdot d_i$. \square

Case	1	2	3	4
i bids b_i	X	X	✓	✓
i bids v_i	X	✓	X	✓

Table 1: Four possible allocation results when i bids truthfully and untruthfully. ✓ denotes i wins the auction, and X denotes i loses the auction.

(3) Auction Truthfulness

Based on the above lemmas, we now prove the main result of this section: the truthfulness of the VERITAS algorithm.

THEOREM 1. *VERITAS spectrum auction is truthful.*

PROOF. Table 1 lists all the possible allocation results when i bids b_i ($b_i \neq v_i$) and v_i . Let u_i^b and u_i^t be i 's utilities when bidding b_i and v_i respectively. By the definition of a truthful auction, we will show that $u_i^t \geq u_i^b$ if $b_i \neq v_i$ in all four cases.

We start from the scenario where i bids higher than its true value, $b_i > v_i$.

- *Case 1:* i loses with both bids. Because VERITAS charges denied bidders by zeros, the utility $u_i^t = u_i^b = 0$, our claim holds.
- *Case 2:* i wins by bidding v_i , yet loses by bidding higher value b_i . From Lemma 1, this cannot happen.
- *Case 3:* Since bidder i is rejected when bidding truthfully, i 's critical neighbor (Lemma 3) must be ranked before i when i bids truthfully. And this is also the neighbor which determines the per-channel clearing price when i bids b_i . It implies that the per-channel clearing price must be greater than (or equal to) i 's true value, which makes $u_i^b \leq 0$. But $u_i^t = 0$, hence $u_i^t \geq u_i^b$.
- *Case 4:* i wins with both bids. By Lemma 4, in both cases VERITAS charges i the same price which is i 's critical value. Hence, $u_i^t = u_i^b$.

Next we consider the scenario where $b_i < v_i$.

- *Case 1:* This case is the same as the Case 1 above.
- *Case 2:* From Lemma 5, when i bids truthfully, its per-channel clearing price is no more than v_i , namely $u_i^t \geq 0$. Since $u_i^b = 0$, we have $u_i^t \geq u_i^b$.
- *Case 3:* i loses by bidding v_i , yet wins by bidding lower value b_i . From Lemma 1, this cannot happen.
- *Case 4:* This is the same as the Case 4 above.

We have shown that when any bidder bids other than its true value, its utility cannot be more than that when it bids its true value. This completes the proof. □

4.3 Flexible Bid Ranking Metrics

The allocation and pricing algorithms in the previous section perform on the bid set that is sorted in descending order of bid b_i . However, it is straightforward to show that VERITAS remains truthful when the bid set is sorted in the descending order of a function $f(b_i)$. The only requirement for $f(b_i)$ is that it is an increasing function of the bid b_i , and not affected by the bids of other bidders. For example, the possible sorting functions are b_i , $\frac{b_i}{|N(i)+1}$ and $b_i \cdot |N(i)|$.

Capable of supporting flexible bid ranking metrics is a key advantage of VERITAS. Auctioneers can design different metrics to tune the VERITAS allocation algorithm towards desired goals. For example, it has been shown that to maximize the sum of winning bids, also known as the *social welfare* [11], the best-known greedy allocation algorithm assigns channel following the descending order of $\frac{b_i}{|N(i)+1}$ [14].

4.4 VERITAS Computational Complexity

We now analyze the running time of VERITAS for a given conflict graph $G = (V, E)$ with n bidders and k channels. First, VERITAS-Alloc takes $O(n \log n)$ time to sort the bids. To allocate channels to bidder i , VERITAS-Alloc needs to examine i 's neighbors for all k channels to find the available lowest-indexed channels. This process takes $2k|E|$ time for n bidders. Therefore, the overall complexity of VERITAS-Alloc is $O(n \log n + k|E|)$. Second, VERITAS-Pricing uses the sorted bids from VERITAS-Alloc and hence its complexity only comes from the process of finding the available lowest-indexed channels, which is $2k|E|$ for each bidder. Therefore, the overall complexity of VERITAS-Pricing is $O(nk|E|)$. Together, the overall complexity of VERITAS with strict requests is $O(n \log n + nk|E|)$.

THEOREM 2. *VERITAS runs in time $O(n \log n + nk|E|)$, where $|E|$ is the number of edges in the conflict graph G , n is the number of bidders, and k is the number of channels auctioned. Because $|E| \leq \frac{n(n-1)}{2}$, VERITAS runs in time less than $O(n^3k)$.*

5. EXTENDING TO OTHER REQUEST FORMATS

In this section, we show that VERITAS can be extended to support different spectrum request formats. In particular, we focus on range-based requests where bidder i requests d_i channels but accepts to obtain any number of channels between 0 and d_i . We also show that VERITAS and Range-VERITAS can be applied to cases where the channels assigned to each bidder are contiguously aligned in frequency.

5.1 Range-VERITAS Auction Design

We start from relaxing the spectrum request format to be range based rather than the strict d_i or 0 channels. We show that VERITAS can be modified slightly to provide the bidders with this flexibility while maintaining truthfulness.

Range-VERITAS-Allocation We modify the VERITAS-Alloc algorithm (Algorithm 1) in step 4 and 5. When the number of available channels for i is less than d_i , we allocate whatever is possible.

Range-VERITAS-Pricing Similarly to the VERITAS-Pricing algorithm, Range-VERITAS charges winning bidders based on the bids of their critical neighbors. The fundamental difference is that now the bidder has multiple (rather than one) critical neighbors because bidding below each critical neighbor will result into the allocation of different number of channels. For each set of additional channels obtained by bidding higher than the last critical neighbor, Range-VERITAS charges with the bid of the last critical neighbor. We list the pricing algorithm in Algorithm 4. Compared to VERITAS-Pricing, the new algorithm differs in step 12, 13, 17 and 18. In step 12 and 13, the bidder is allocated

Algorithm 3 Range-VERITAS-Alloc(B, D, G)

```
1:  $B'$  =sorted  $B$ 
2: while  $B' \neq \emptyset$  do
3:    $i$  =TOP( $B'$ );
4:   if  $Distinct(N(i)) \leq k$  then
5:     Assign( $i, \min(d_i, k - Distinct(N(i)))$ )
6:   end if
7:    $B' = B' \setminus \{b_i\}$ 
8: end while
```

Algorithm 4 Range-VERITAS-Pricing(B, D, G, i)

```
1:  $B'$  =sorted  $B$ 
2:  $owned\_ch = \sum_{j=1}^k a_{ij}$ 
3: if  $owned\_ch = 0$  then
4:    $p_i = 0$ 
5:   return
6: end if
7:  $p_i = 0$ 
8:  $avail\_ch = \{c_1, c_2, \dots, c_k\}$ 
9:  $B'' = B' \setminus \{b_i\}$ 
10: while ( $B'' \neq \emptyset$ ) AND ( $owned\_ch > 0$ ) do
11:    $q$  =TOP( $B''$ )
12:   if  $Distinct(N(q)) < k$  then
13:     Assign( $q, \min(k - Distinct(N(q)), d_q)$ )
14:     if  $q \in N(i)$  then
15:        $avail\_ch = avail\_ch \setminus$  channels allocated to  $q$ 
16:       if  $|avail\_ch| < owned\_ch$  then
17:          $p_i = p_i + b_q \cdot (owned\_ch - |avail\_ch|)$ 
18:          $owned\_ch = |avail\_ch|$ 
19:       end if
20:     end if
21:   end if
22:    $B'' = B'' \setminus \{b_q\}$ 
23: end while
```

with channels if there is any available; in step 18, we update $owned_ch$ to the number of currently available channels so that the clearing price for each bundle of channels is accumulated in step 17. The total clearing price for i is the sum of prices charged for all of its assigned channels.

Toy Example We use the example in Figure 5 to explain how Range-VERITAS charges winning bidders. Assume $k = 3$ and $d_i = 2$. Figure 5 (left) shows the allocation result when every bidder is truthful. To compute the clearing price, say for bidder e , we consider the bidder list $\langle a, b, c, d \rangle$. Using Range-VERITAS-Alloc, e 's available channels are not affected until c is allocated with 2 channels $\{1, 2\}$. Then there is only 1 available channel for e , which is less than 2, the number of channels e obtained from the auction. This means by bidding below c , e can obtain at most 1 channel. Hence for that $(2 - 1)$ channel, e is charged by c 's bid 2. When e 's neighbor d is allocated with channel $\{3\}$, there is no available channel for e , then for that left incremental 1 channel, e is charged by d 's bid 1. So the total price e needs to pay for $\{1, 2\}$ is $2 + 1 = 3$, and e 's utility is $3 * 2 - 3 = 3$.

Figure 5 (right) shows the auction result if c lies by bidding 4 instead of the true value 2. Now let us look at the clearing price for c . Consider the bidder list $\langle a, b, e, d \rangle$. Using Range-VERITAS-Alloc, a first gets channels $\{1, 2\}$ and b gets channel $\{3\}$. Since b is c 's neighbor, now c has 2 available channels. Next e is allocated with 2 channels




Figure 5: An example showing how Range-VERITAS-Pricing works. When bidder c raises its bid (right part), it obtains a negative utility.

$\{1, 2\}$, so the number of c 's available channels reduces to 0, which is less than the channels c obtained from auction. Therefore for these $(2 - 0)$ channels, c is charged by e 's bid 3 per channel, which means c 's utility is $(2 - 3) * 2 < 0$. That is, c cannot obtain a higher utility by bidding untruthfully.

5.2 Range-VERITAS Truthfulness

Consider two different bids b_i^1 and b_i^2 , $b_i^1 < b_i^2$. Let d_i^1 and d_i^2 denote the number of channels i obtains in each case; p_i^1 and p_i^2 represent the total clearing price at i ; $pos(b_i^1)$ and $pos(b_i^2)$ be the positions of the bidder i in the sorted bid lists using the two bids.

LEMMA 6. When $b_i^1 < b_i^2$, the following statements hold:

1. $d_i^1 \leq d_i^2$.
2. The amount charged for d_i^1 ($d_i^2 = d_i^1 + \delta$, where $\delta \geq 0$) channels is the same for both bids.
3. When the bid is b_i^2 , for the exactly $(d_i^2 - d_i^1)$ channels, the per-channel clearing price is greater than b_i^1 .
4. When the bid is b_i^2 , for the exactly $(d_i^2 - d_i^1)$ channels, the per-channel clearing price is no more than b_i^2 .

PROOF. For statement 1, the arguments are the same as those used in Lemma 1 – Because all the bidders before $pos(b_i^2)$ are the same in both the sorted bid lists (Figure 4), the higher bid (b_i^2) obtains at least the same number of channels as that of the lower bid (b_i^1).

For statement 2, consider the two sorted lists of bids in Figure 4. Since i gets d_i^1 channels at position $pos(b_i^1)$, the critical neighbors for the d_i^1 channels must be located after position $pos(b_i^1)$. By using these critical neighbors' bids, Range-VERITAS calculates the prices for these d_i^1 channels. Since the list of bidders after $pos(b_i^1)$ does not change, both cases will charge the same amounts for the d_i^1 channels.

Statement 3 follows the argument that only d_i^1 channels are available at position $pos(b_i^1)$. Range-VERITAS will use one or more neighbors located before $pos(b_i^1)$ to determine the prices for the additional $(d_i^2 - d_i^1)$ channels. Since all bids before $pos(b_i^1)$ are greater than b_i^1 , the per-channel price for the $(d_i^2 - d_i^1)$ channels is greater than b_i^1 .

The argument for statement 4 comes from those of Lemma 5. For each bundle of channels i obtains, the price is determined by the neighbor below whom i will lose that bundle of channels. Since i does get additional $(d_i^2 - d_i^1)$ channels at $pos(b_i^2)$, any bid that pushes i before $pos(b_i^2)$ will also let i get them. Hence the price for $(d_i^2 - d_i^1)$ channels cannot

be determined by the neighbors before $pos(b_i^2)$, namely the per-channel price for the $(d_i^2 - d_i^1)$ channels cannot be more than b_i^2 . \square

Using the above statements, we now prove that Range-VERITAS is truthful. Again, we consider all the possible cases listed in Table 1. Note that \surd represents the case when the bidder obtains one or more channels.

THEOREM 3. *For range requests, Range-VERITAS spectrum auction is truthful.*

PROOF. We start from the case when i bids higher than its true value and show $u_i^t \geq u_i^b$ if $b_i > v_i$.

- *Case 1-3:* The arguments are similar as those in the corresponding cases in Theorem 1.
- *Case 4:* i wins with b_i or v_i . Let d and $(d + \Delta d)$ represent the number of channels assigned to i when bidding v_i and b_i . From Lemma 6, $\Delta d \geq 0$. Let p and $(p + \Delta p)$ represent the clearing prices in each case. From Lemma 6, part 2, we know that the d channels are charged the same, then we have $u_i^t = (v_i d - p)$, and $u_i^b = (v_i d - p) + (v_i \Delta d - \Delta p)$. From Lemma 6, part 3, the per-channel price for Δd is larger than v_i , namely $\Delta p > v_i \Delta d$. Therefore $u_i^b = (v_i d - p) + (v_i \Delta d - \Delta p) < (v_i d - p) = u_i^t$.

Next we show that when a bidder bids lower than its true value $b_i < v_i$, it cannot increase its utility.

- *Case 1-3:* They are similar as those in Theorem 1.
- *Case 4:* i wins with b_i or v_i . Let d and $(d + \Delta d)$ represent the number of channels assigned to i when bidding b_i and v_i . From Lemma 6, $\Delta d \geq 0$. Let p and $(p + \Delta p)$ represent the clearing prices in each case. Again, from Lemma 6, part 2, we know that the d channels are charged the same, then we have $u_i^b = (v_i d - p)$, and $u_i^t = (v_i d - p) + (v_i \Delta d - \Delta p)$. From Lemma 6, part 4, the per-channel price for Δd is no more than v_i , namely $\Delta p \leq v_i \Delta d$. Therefore $u_i^t = (v_i d - p) + (v_i \Delta d - \Delta p) \geq (v_i d - p) = u_i^b$.

From the above, we have shown that for all possible cases, no bidder can improve its utility by bidding untruthfully. \square

5.3 Contiguous Spectrum Requests

VERITAS and Range-VERITAS can be easily extended to scenarios where bidders require their allocated channels to be contiguously aligned in frequency. The pricing mechanisms remain the same, while in the allocation processes, we check whether d_i or $x \leq d_i$ contiguously aligned channels are available. We can show that the modified auction algorithms are still truthful. The proof follows on similar lines as Theorem 1. In the interest of space, we omit the proof.

THEOREM 4. *Contiguous-VERITAS algorithm is truthful.*

THEOREM 5. *For contiguous range based bids, Contiguous-Range-VERITAS algorithm is truthful.*

6. EXPERIMENTAL RESULTS

In this section, we perform experiments to evaluate the performance of VERITAS. First, we explore the unique property of truthful auctions by comparing VERITAS to a revenue-maximizing yet non-truthful auction. Second, we examine

the *efficiency* of VERITAS by comparing it to the simple truthful auction described in Section 3.3. We also compare VERITAS's greedy allocation to the best-known greedy mechanisms in spectrum allocation. Finally, we investigate the impacts of different ranking metrics and allocation patterns to demonstrate VERITAS's *flexibility*.

We assume a single auctioneer that handles bidders in a large geographic area. We randomly deploy bidders in a square 1×1 area, and apply a distance-based interference model to produce the corresponding conflict graph. Any two bidders within 0.1 distance will conflict with each other and hence cannot be allocated with the same channels. We assume that each bidder's true valuation (and hence its bid) is uniformly distributed over $(0, 1]$. By default, each bidder requests one channel. The results are averaged over 5 random seeds.

We use the following three performance metrics.

- *Revenue:* The sum of charges to all the winners.
- *Spectrum Utilization:* The sum of allocated channels of all the winning bidders.
- *User Satisfactory:* The percentage of winning bidders. When each bidder demands only one channel, it is the ratio of the spectrum utilization to the total number of bidders. Otherwise, it is the percentage of bidders who are allocated with one or more channels.

6.1 Truthful vs Non-Truthful Auctions

To explore the property of truthfulness, we compare VERITAS with a non-truthful revenue-maximizing auction. The non-truthful scheme chooses winners to maximize the revenue and charges winners by their actual bids. When each bidder requires only one channel, the clearing problem in non-truthful auction reduces into a maximum weighted independent set problem. We use the well-known solutions from [14]. Given a set of bids, the non-truthful auction always produces higher revenue compared to VERITAS since VERITAS charges winners less than their actual bids. However, this comparison is unfair because bidders in non-truthful auctions have no incentive to bid their true values. Nevertheless, we plot the results of both auctions based on the same set of bids to illustrate the trends of spectrum utilization and revenue as the number of bidders and the available channels increase.

In Figure 6, we plot the spectrum utilization and revenue for both the non-truthful auction ((a)-(c)) and VERITAS ((d)-(f)). Both systems perform similarly in terms of spectrum utilizations, which grow with the number of channels auctioned and the number of players, and eventually saturate. On the other hand, the two systems behave differently in terms of the revenue as the number of channels increases. The revenue of the non-truthful auction keeps increasing, while the revenue of VERITAS starts to decline once the number of channels auctioned exceeds a value.

This significant difference comes from the fact that the two auctions have fundamentally different charging mechanisms: the non-truthful auction charges winners by their actual bids while the truthful auction (and VERITAS) charges winners by the bids of their critical neighbors. In the non-truthful auction, increasing the number of channels also increases the number of winners, and hence the revenue which is the sum of winning bids. In VERITAS, although the number of winners increases, the charge to individual winner decreases as

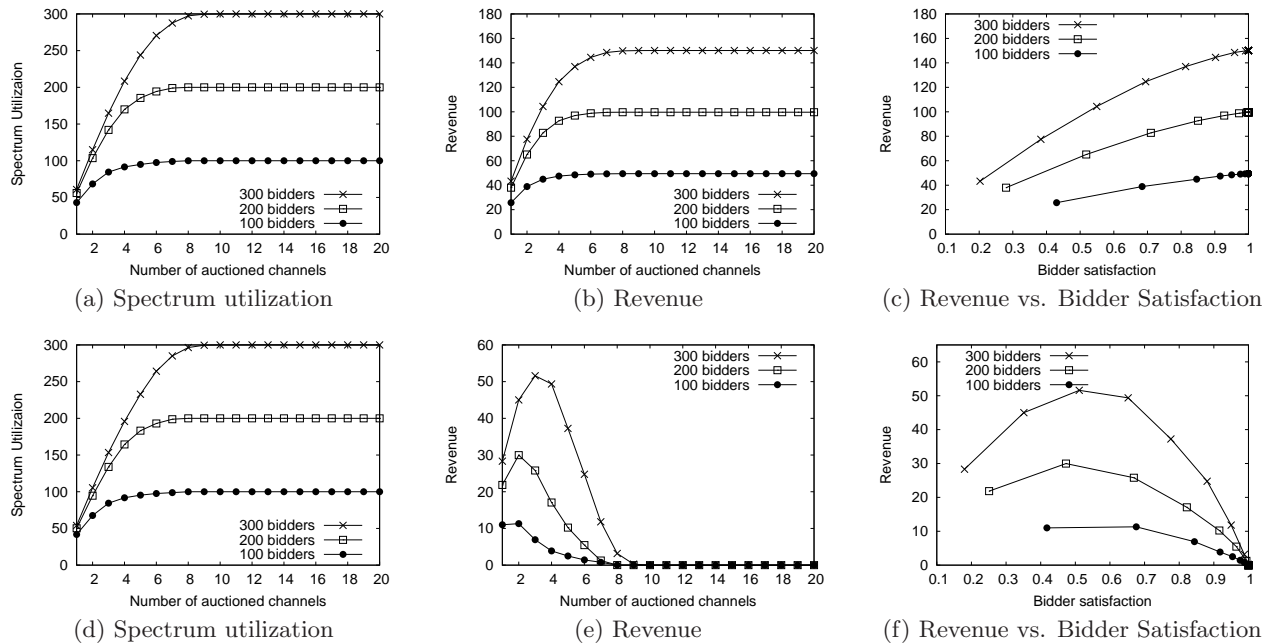


Figure 6: Comparing VERITAS to the non-truthful revenue-maximizing spectrum auction. Top 3 figures assume the non-truthful spectrum auction and bottom 3 are for VERITAS.

the pool of losing bidders shrinks. For a better illustration, Figure 6(c) and (f) plot the revenues against the user satisfactions. We see that for VERITAS, there exists an optimal user satisfactory level which balances the number of winners and their charges to maximize the revenue.

Motivated by this interesting phenomenon, we propose to let the auctioneer determine the number of channels auctioned in order to maximize its revenue. In Section 7, we will discuss analytically how to choose the number of channels auctioned given the information of the conflict graph and the bid distribution.

6.2 Simple vs Sophisticated Truthful Auctions

Next, we compare VERITAS with the simple truthful scheme in Section 3.3. The key distinction between them is that VERITAS seeks to improve the efficiency in spectrum usage, while the simple auction design sacrifices spectrum utilization to convert the problem into the conventional auction setting. Given the interesting revenue-satisfaction tradeoff in Figure 6(f), a natural question is “*Can VERITAS’s improved efficiency in spectrum usage also lead to higher revenue?*”

To answer this question, we compare VERITAS and the simple auction in two scenarios. First, we fix the number of channels auctioned and vary the number of bidders. Second, we fix the number of bidders but vary the number of channels auctioned. These experiments allow us to understand both approaches under different levels of resource contention.

1) *Varying the number of bidders* – From Figure 7(a)-(b), we see that the simple truthful auction performs poorly in terms of spectrum utilization because it sacrifices spectrum reuse to maintain truthfulness. However, when the number of bidders is small, the simple design produces higher revenue than VERITAS. This is because when less than 200

bidders compete for 8 channels, more than 90% of them succeed in VERITAS but less than 50% succeed in the simple design. For a fair comparison, Figure 7(c) compares both approaches in terms of the tradeoff between revenue and user satisfaction. Given the same bidder satisfaction, VERITAS increases the revenue significantly by up to 200%.

2) *Varying the number of channels auctioned* – Similarly, Figure 8(a) shows that VERITAS improves spectrum utilization significantly. From Figure 8(b), we observe that although the maximal revenues of both auction systems are similar, VERITAS achieves the maximum using 3 channels while the simple design requires 10 channels. For a fair comparison, we plot the revenue per channel as a function of the bidder satisfactory rate by varying the number of channels auctioned (Figure 8(c)). Again, VERITAS significantly outperforms the simple design.

6.3 Efficiency of VERITAS

In Section 6.2, we have shown that improving allocation efficiency is a critical requirement for truthful auctions. Next, we evaluate VERITAS’s efficiency by comparing its allocation algorithms to two best-known greedy algorithms that maximize the spectrum utilization and the social welfare (sum of winning bids). In particular, we compare:

- Greedy - $1/(X + 1)$ [14]: The best-known greedy allocation algorithm that maximizes the spectrum utilization. This algorithm allocates channels sequentially. Each time it sorts bidders by ranking $1/(X + 1)$ where $X = |N(i)|$ is i ’s conflict degree, and assigns channels to the highest-ranked bidder. Note that this algorithm updates X s after each allocation, while VERITAS keeps the same order to maintain truthfulness.
- Greedy - $b/(X + 1)$ [14]: The best-known greedy al-




Figure 7: Comparing VERITAS and the simple truthful design by auctioning 8 channels to 10–500 bidders.




Figure 8: Comparing VERITAS and the simple truthful design by auctioning 1–30 channels to 300 bidders.

location that maximizes the sum of winning bids (or *social welfare* [11]) by using a ranking of $b/(X + 1)$. This algorithm also updates X s and reorders bidders after each allocation.

- VERITAS - b : VERITAS with a ranking metric of b .
- VERITAS - $b/(X + 1)$: VERITAS with a ranking metric of $b/(X + 1)$. By considering the conflicting degree X , this ranking metric intends to maximize the social welfare. Unlike Greedy- $b/(X + 1)$, it does not reorder bidders after each allocation in order to maintain truthfulness.
- VERITAS - $1/(X + 1)$: As a reference, VERITAS with a ranking metric of $1/(X + 1)$.

The comparisons are fair since all these algorithms are of polynomial complexity. Both greedy algorithms are within proven distances to the optimal solution (NP-complete), and serve as good benchmarks to evaluate VERITAS.

Random Topologies Figure 9(a)-(c) compare the greedy and VERITAS algorithms in terms of the spectrum utilization, the sum of winning bids, and the revenue, assuming each bidder requests one channel. As expected, the $1/(X + 1)$ metric maximizes the spectrum utilization while the $b/(X + 1)$ metric maximizes the sum of winning bids. More importantly, VERITAS achieves similar performance (<5% degradation) as those greedy algorithms. The performance degradation in VERITAS comes from the fact that the greedy allocation algorithms update X value after each allocation while VERITAS must use the original X to enforce truthfulness. Finally, with a uniformly random conflict

graph, b and $b/(X + 1)$ metrics in VERITAS perform similarly.

Clustered Topologies To create clustered topologies, we initially distribute 60 bidders in a 0.5×0.5 area, and then increase bidders up to 300 by adding 60 bidders in the center each time. Figure 9(d)-(f) compare VERITAS to the greedy algorithms. As the level of clustering increases, the use of the ranking metric $b/(X + 1)$ gradually outperforms b by considering the impact of conflict degree.

The above results clearly demonstrate VERITAS’s allocation efficiency. We note that VERITAS also allows the auctioneer to optimize customized utility functions other than the revenue. For example, to maximize the social welfare, VERITAS can choose $b/(X + 1)$ as the ranking metric.

6.4 Supporting Different Demand Patterns

In addition to providing flexibility to the auctioneer, VERITAS also offers flexible demand formats to the bidders. Next, we examine how these different demand patterns impact the auction results. We compare: strict requests ($0/d_i$) where d_i represents bidder i ’s demand, range requests $[0, d_i]$, and continuous strict requests ($0/\text{Cont } d_i$).

Figure 10(a)-(b) compare the spectrum utilization and revenue of the three patterns when auctioning channels to 200 bidders with $d_i = 10$. We observe that the use of range requests leads to the highest spectrum utilization and maximizes revenue using less channels. In Figure 10(c), we plot revenue as a function of the number of channels auctioned using range requests. The results show that the optimal number of channels auctioned depends heavily on d_i .

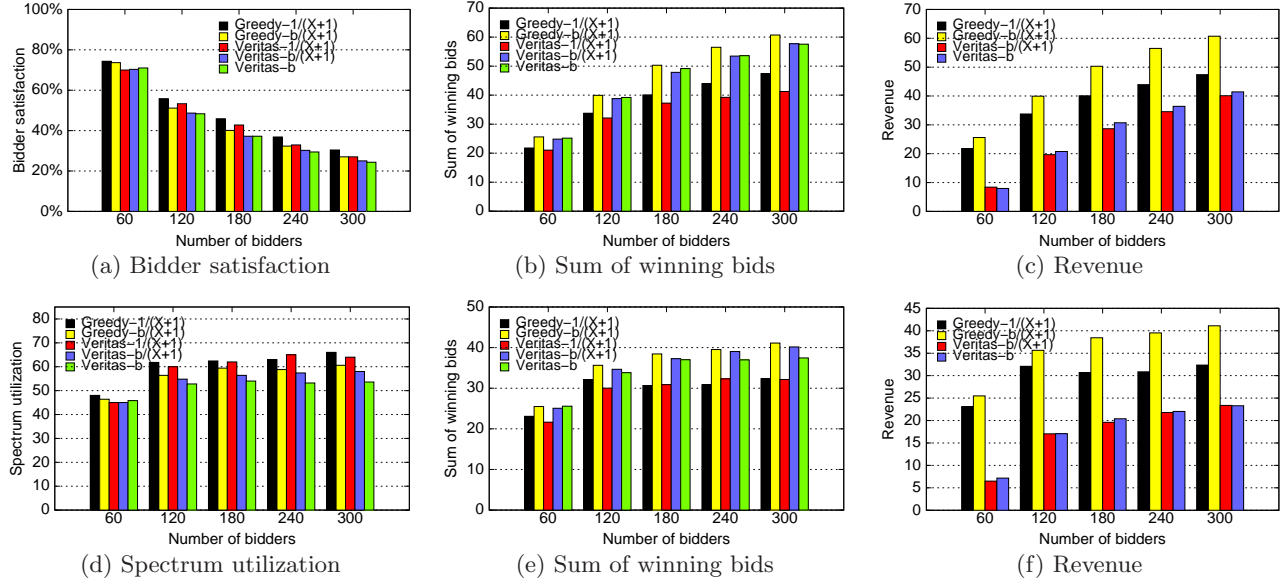


Figure 9: Comparing the performance of VERITAS to untruthful greedy allocations with different ranking metrics. Top 3 figures assume random conflict graphs and bottom 3 figures assume clustered conflict graphs.




Figure 10: VERITAS under different demand/allocation patterns. (a)-(b) $d_i=10$, (c) varying d_i : 3,5,10.

7. MAXIMIZING REVENUE BY OPTIMIZING THE NUMBER OF CHANNELS AUCTIONED

As discussed in Section 6.1, the auctioneer can maximize its revenue by choosing the number of channels auctioned k_{opt} . The choice of k_{opt} should be made before the actual auction to maintain truthfulness. While this optimization can be highly complex due to its dependency on the network topology, we provide an analytical framework to derive k_{opt} assuming the conflict topology is all-connected (clique) and the bids are independent and identically-distributed. In this case, VERITAS's pricing reduces to k -position pricing [15]. Finally, we use simulations to evaluate its performance when it is applied to non-clique conflict graphs.

Consider n bidders competing for k channels and each bidder demands one channel. Let $b_1 \geq b_2 \geq \dots \geq b_n$. In this case, VERITAS assigns channels to the top k ranked bidders and charge each of them with b_{k+1} . Hence the total revenue is $k \cdot b_{k+1}$, which is a random variable. We assume that

each bidder's bid b_i ($i = 1..n$) follows the same independent distribution with the probability density function $p(x)$ and cumulative distribution function $F(x)$. Our problem is to find the optimal k that maximizes the expected revenue:

$$k_{opt} = \underset{k=1..n}{\operatorname{argmax}} E(k \cdot b_{k+1}) = \underset{k=1..n}{\operatorname{argmax}} \int k p_{b_{k+1}}(x) \cdot x dx \quad (1)$$

where $p_{b_{k+1}}(x)$ is the probability density function for b_{k+1} :

$$p_{b_{k+1}}(x) = \frac{n!}{(n-k-1)!k!} F(x)^{n-k-1} \cdot (1-F(x))^k \cdot p(x) \quad (2)$$

When the bid follows uniform distribution in the range of $(0, 1]$, that is $p(x) = 1$ and $F(x) = x$, then we can derive

$$E(k \cdot b_{k+1}) = \frac{n!}{(n-k-1)!(k-1)!} \int_0^1 x^{n-k} (1-x)^k dx \quad (3)$$

and

$$k_{opt}(n) = \begin{cases} \frac{n}{2}, & n \text{ is even} \\ \frac{n-1}{2}, \frac{n+1}{2}, & n \text{ is odd} \end{cases} \quad (4)$$

In a random network, we propose to approximate k_{opt} by the average conflict degree n plus one. Figure 11 compares

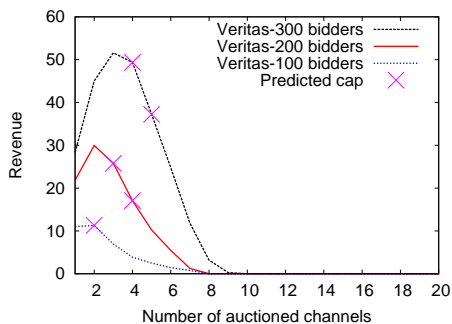


Figure 11: Theoretical prediction of the optimal number of channels to be auctioned.

our predicted k_{opt} and the actual performance of VERITAS in non-clique conflict graphs. We vary the number of bidders from 100 to 300 in the same deployment area, mapping to the average conflict degrees of 3, 6, 8, and hence n of 4, 7, 9 in (4). We observe that the proposed predictions are very close to the actual optimal values. As the conflict degree increases, our prediction requires more channels because it conservatively assumes that local interference graphs are all-connected. Finally, we note that the optimizations for other types of network topologies are much harder, and we seek to address them in a future study.

8. RELATED WORK

Auctions have been used by FCC to allocate scarce spectrum resource [3]. Huge body of works follow on the designs of wireless spectrum auctions in different scenarios. These include transmit power auctions [6] where bidders use the same spectrum band but bid transmit power to minimize the interference, and spectrum band auctions [1, 4, 7, 13] where bidders obtain different spectrum channels to minimize the interference. However, none of these addresses the problem of truthfulness/strategy-proofness.

Truthfulness/strategy-proofness is a crucial property from the economic perspective. It has been applied to other scenarios beyond traditional auctions such as the multi-cast routing [16]. The notion of truthful bidding in sealed-bid auctions was first brought out by Vickery in his seminal paper [15], which introduced the secondary pricing for one-item auctions. This work was then generalized to multiple-unit combinatorial auctions by Clarke [2] and Groves [5], resulting into the famous Vickery-Clarke-Groves (VCG) mechanism. To achieve truthfulness, this celebrated scheme requires the optimal welfare-maximizing allocation, which unfortunately is a NP-complete problem in general multi-unit auctions. Since then, there have been some works on truthful mechanism design using greedy algorithms or linear programming for specific classes of bidders in specific problem domains [10] [11]. We point the interested readers to the survey in [12] for more details.

As discussed in Section 3, these existing solutions, when applied to spectrum auctions, either lose the truthfulness, require exponential complexity, or result in significant degradation of spectrum utilization. We build VERITAS using the observations from existing proposals, and achieve truthfulness in a spectrum-efficient and computationally-efficient

manner. VERITAS also allows the auctioneer to tailor its allocation mechanism towards different economic or spectrum objectives, and provides bidders with diverse bidding formats.

9. CONCLUSIONS

In this work, we propose VERITAS, a truthful and efficient dynamic spectrum auction system to serve many small players. Like an eBay marketplace, VERITAS allows wireless users to obtain and pay for the spectrum based on their demands, and enables spectrum owners to maximize their revenues by assigning spectrum to the bidders who truly value it the most. We prove analytically VERITAS's truthfulness and computational efficiency. We also show that VERITAS is highly efficient and flexible, and can be easily reconfigured to suit multiple needs of the auctioneer and bidders.

Acknowledgement

We thank the anonymous reviewers for their insightful suggestions. We gratefully acknowledge the support from NSF through grant CNS-0721961.

10. REFERENCES

- [1] BUDDHIKOT, M., AND RYAN, K. Spectrum management in coordinated dynamic spectrum access based cellular networks. In *Proc. of IEEE DySPAN* (2005).
- [2] CLARKE, E. Multipart pricing of public goods. *Public Choice XI* (1971), 17–33.
- [3] CRAMTON, P. Spectrum auctions. *Handbook of Telecommunications Economics* (2002), 605–639.
- [4] GANDHI, S., BURAGOHAJ, C., CAO, L., ZHENG, H., AND SURI, S. A general framework for wireless spectrum auctions. In *Proc. of IEEE DySPAN* (2007).
- [5] GROVES, T. Incentive in terms. *Econometrica* 41 (1973), 617–631.
- [6] HUANG, J., BERRY, R., AND HONIG, M. Auction mechanisms for distributed spectrum sharing. In *Proc. of 42nd Allerton Conference* (2004).
- [7] ILERI, O., SAMARDZIJA, D., AND MANDAYAM, N. B. Demand responsive pricing and competitive spectrum allocation via a spectrum server. In *Proc. of IEEE DySPAN* (2005).
- [8] JAIN, K., PADHYE, J., PADMANABHAN, V. N., AND QIU, L. Impact of interference on multi-hop wireless network performance. In *Proc. of MobiCom* (2003).
- [9] KOTHARI, A., PARKE, D. C., AND SURI, S. Approximately-strategyproof and tractable multi-unit auctions. In *Proc. of the ACM conference on Electronic Commerce* (2003), pp. 166–175.
- [10] LEHMANN, D., O'CALLAGHAN, L., AND SHOHAM, Y. Truth revelation in approximately efficient combinatorial auctions. *J. ACM* 49, 5 (2002), 577–602.
- [11] MU'ALEM, A., AND NISAN, N. Truthful approximation mechanisms for restricted combinatorial auctions: extended abstract. In *Eighteenth national conference on Artificial intelligence* (2002), pp. 379–384.
- [12] NISAN, N., ROUGHGARDEN, T., TARDOS, E., AND VAZIRANI, V. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [13] RYAN, K., ARAVANTINOS, E., AND BUDDHIKOT, M. A new pricing model for next generation spectrum access. In *Proc. of TAPAS* (2006).
- [14] SAKAI, S., TOGASAKI, M., AND YAMAZAKI, K. A note on greedy algorithms for the maximum weighted independent set problem. *Discrete Appl. Math.* 126, 2-3 (2003), 313–322.
- [15] VICKERY, W. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance* 16 (1961), 8–37.
- [16] WANG, W., LI, X., AND WANG, Y. Truthful multicast routing in selfish wireless networks. In *Proc. of MobiCom* (2004).