Breaking Bidder Collusion in Large-Scale Spectrum Auctions

Xia Zhou
Department of Computer Science
University of California, Santa Barbara
CA 93106, USA
xiazhou@cs.ucsb.edu

Haitao Zheng
Department of Computer Science
University of California, Santa Barbara
CA 93106, USA
htzheng@cs.ucsb.edu

ABSTRACT

Dynamic spectrum auction is an effective solution to provide spectrum on-demand to many small wireless networks. As the number of participants grows, bidder collusion becomes a serious threat. In this paper, we study bidder collusion in large-scale spectrum auctions, investigating its impact on auction outcomes. We found that the nature of the complex interference constraints among bidders provides a fertile breeding ground for colluders, causing significant damage in auction efficiency and revenue. In particular, collusion group of small size plays a dominant role since it is easy to form and hard to be detected.

We propose Athena, a new collusion-resistant auction framework for large-scale dynamic spectrum auction. Athena implements a soft collusion resistance, allowing the auctioneer to exploit the tradeoff between the level of collusion resistance and the cost of achieving such level of resistance. Unlike existing solutions, Athena enables spectrum reuse across bidders, achieves soft collusion resistance against any form of collusive bidding strategy, maintains provable revenue guarantee, and does so with polynomial-time complexity. To provide a comprehensive evaluation, we first analytically prove Athena's collusion resistance and revenue guarantee (under any bids), and then experimentally verify our analytical conclusions using empirical bid distributions.

Categories and Subject Descriptors

C.2 [**Computer-Communication Networks**]: Network Architecture and Design

General Terms

Algorithms, Design

Keywords

Dynamic spectrum auctions, collusion-resistance, cognitive radio networks

MobiHoc'10, September 20–24, 2010, Chicago, Illinois, USA.

1. INTRODUCTION

Small wireless networks are blooming across the globe. In New York City alone, the number of WiFi networks totaled more than 14,000 in 2002 and it continues to grow [7]. Existing deployments, however, use unlicensed spectrum bands for low-cost entry and rapid growth. Now they suffer excessive interference and poor performance due to the aggressive deployment and unprotected spectrum usage. Soon, these networks will stop growing and even collapse due to spectrum shortage, despite their increasing popularity.

But the future is bright. Recent changes in FCC spectrum policy push for secondary spectrum markets where small wireless networks can obtain additional spectrum. Several proposals [3, 9, 15, 25] have shown that dynamic spectrum auction is an effective platform for secondary markets. These new auctions target small wireless networks, allowing them to bid for spectrum by their short-term local usages. Properly designed auctions will assign protected spectrum to networks producing the best economic outcomes.

To be successful, an auction must be resilient to selfish bidders. Strategic bidders, individually or in groups, seek to game the system by rigging their bids to manipulate auction outcomes and to improve their own utilities. To make the best use of the spectrum, an auction must discourage bidders from cheating and instead encourage them to reveal their true valuations of the spectrum to the auctioneer. In this context, prior work has developed truthful dynamic spectrum auctions to discourage individual cheating [17, 25, 26]. Truthful auctions ensure that no bidder, individually, can improve its utility by bidding other than its true valuation. Thus a rational strategy is to bid its true valuation.

Truthful auctions, however, become ineffective when bidders collude, *i.e.* when bidders strategically form collusion groups and rig their bids together to manipulate auction results [2, 6]. Designed to address individual cheating, truthful auctions cannot prevent collusion groups from improving their group utilities. In fact, collusion has appeared in several past commercial auctions, and has caused significantly lower auction revenue and unfair resource distribution [2, 5, 6, 8, 12]. Interestingly, existing measurements also show that collusion groups were in general tacit and small in size, because they are easy to form and hard to be detected.

In this paper, we study bidder collusion in the context of emerging large-scale dynamic spectrum auctions. Our work differs from prior work [2, 5, 6, 8, 12] because these new spectrum auctions must consider spectrum reusability. Unlike books or paintings, radio spectrum is reusable across bidders. The competition among bidders is now defined by a large set of complex interference constraints. The nature of these constraints not only provides a fertile breeding ground for collusion, but also complicates the auction design [25].

To understand and address bidder collusion in spectrum auctions, our study seeks to answer two key questions:

(1) Is bidder collusion, particularly small-size collusion, a big threat to dynamic spectrum auctions?

(2) If so, how can one design auctions tactically to deal with them, and what is the cost for adding such robustness?

To examine the impact of bidder collusion, we start from experimenting on a state-of-the-art truthful spectrum auction design [25]. We show that a simple collusion pattern involving only two bidders can easily improve the group's utility, no matter how others bid. Now bidders have incentive to collude and cheat, many collusion groups will form, degrading auction revenues by up to 50%.

To resist collusion, we take a proactive prevention approach, because uncovering collusion groups is hard due to its tacit nature and the complex auction structure. Specifically, we redesign the auction rules to diminish the gain of collusion groups, leaving bidders little or no incentive to collude. In this context, it is proven that the *only* solution to suppress collusion of any group size and any type of collusive strategy is a trivial posted-price design, which unfortunately leads to unbounded loss in revenue [11]. To reduce the cost, we consider a "soft" approach that targets small-size collusion of any type and suppresses the gain of collusion probabilistically [11, 20]. This approach is shown to be highly effective and cost-efficient in conventional large-scale auctions. Existing designs [11,20], however, target conventional auctions without reusability. Directly applying or extending them to dynamic spectrum auctions either breaks the collusion resistance or creates excessive interference.

The Athena Spectrum Auction. We propose Athena, a new framework for collusion resistant spectrum auction design. Like [11,20], Athena applies soft collusion resistance to address collusion groups of small-sizes, while minimizing the cost in auction revenue for achieving such resilience. Different from [11,20], Athena operates under the complex bidder interference constraints and exploits spectrum reusability to service a large number of small networks. While recent work [23, 24] also considers spectrum reuse, it only handles three specific types of small-size collusion and requires exponential complexity to ensure the collusion resistance. Different from [23, 24], Athena handles all types of small-size collusion, and achieves both collusion resistance and revenue guarantee using computationally-efficient algorithms, making it a low-cost and deployable solution for large-scale spectrum auctions.

The key challenge in our design is that bidders now can exploit the localized interference constraints to build collusion that conventional designs cannot handle. Athena addresses this challenge using a 3-stage "Divide, Conquer, and Combine" process. Using an allocation algorithm, Athena first divides bidders into segments, removing interference constraints within each segment. Athena then suppresses "intra-segment collusion" within each segment using a classical randomized mechanism to select potential winners within the segment. Finally, and most importantly, Athena combines results from all the segments to judiciously select winning segments, preventing "inter-segment collusion" (when bidders across different segments collude) from affecting auction outcomes. As a result, bidders in Athena auction have little incentive to collude, because, regardless of their collusive bidding strategy, neither cheating within each segment nor cheating across multiple segments will produce much gain.

Integrated Evaluation. When evaluating Athena, we encountered the issue on reasoning how bidders will evaluate the spectrum and bid in the auction, which is also an open problem in the field of economics [18]. We approach the solution of this problem by combining theoretical analysis with experimental evaluation. We first analytically prove Athena's properties, and then examine Athena experimentally under illustrative bid distributions. Using the following two case studies, we evaluate Athena by its effectiveness in resisting collusion and the cost of achieving such robustness.

- *Case Study I*: To examine Athena's effectiveness in resisting collusion, we first analytically prove that, for any bids, Athena achieves the (t, p)-truthfulness of [11, 20]. It ensures that with a probability p or higher, no colluding group of size t or less can gain any benefit. The same (t, p)-truthfulness holds when multiple collusion groups are present, as long as each group is of size t or less. The choice of (t, p) determines the level of collusion resistance. To verify this effect experimentally, we identify some collusion patterns that are effective in truthful auctions, and examine their effectiveness under Athena. Our results confirm that these collusion groups cannot improve their group utilities. Together, the analytical proof and the illustrative examples offer solid verifications on Athena's collusion resistance.

- *Case Study II*: To understand the cost of Athena's collusion resistance, we consider a different scenario. Aware of Athena's collusion resistance, bidders now have no incentive to cheat but bid by their true valuations. In this "ideal" scenario, we compare Athena's revenue to that of the truthful auction without any collusion resistance. The difference in revenue indicates the necessary cost for adding such collusion resistance. We first prove that for any bids (or valuations), Athena's revenue is within a bounded distance from the optimal. Next, using several valuation distributions, we examine the revenue loss experimentally in greater detail. Results show that the cost scales gracefully with the level of collusion resistance (t, p), but remains small <15% for $(t = 4, p = 0.8)$.

Summary of Contributions. Targeting emerging large-scale dynamic spectrum auctions, our work makes two key contributions. First, we show that existing spectrum auction designs are highly vulnerable to bidder collusion, especially small-size collusion that is easy to form and hard to be detected. Although each small collusion group leads to minor impact on auction outcome, many collusion groups, together, can significantly damage auction performance. Second, we present Athena, a new collusion-resistant spectrum auction design. Athena effectively resists small-size collusion by diminishing its gain statistically, leaving bidders little or no incentive to collude. Different from prior solutions, Athena not only restores collusion resistance, but also enables spectrum reuse, achieves revenue guarantee, and does so with polynomial-time complexity. To our best knowledge, Athena is the first large-scale spectrum auction design that achieves such general and cost-efficient collusion resistance.

2. PRELIMINARIES

As background, in this section we introduce the emerging large-scale dynamic spectrum auctions, and provide a brief overview of bidder collusion in conventional auctions.

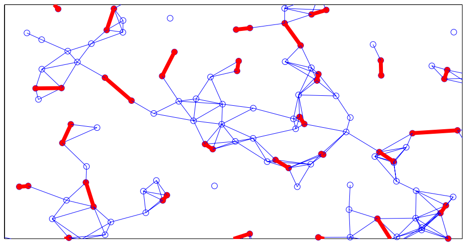

Figure 1: A sample network area that describes bidder conflict conditions. Two bidders conflict with each other if they are directly connected. WCN groups are marked in red. W and C are conflicting neighbors.

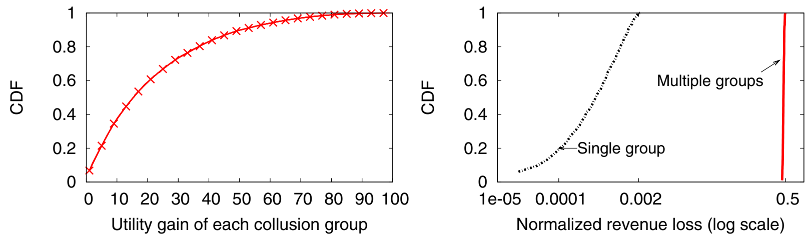

(a) Gain in group utility (b) Revenue loss due to collusion

Figure 2: WCN collusion is effective in large-scale spectrum auction like VERITAS [25]. (a) It can effectively improve group utility, giving bidders incentive to collude. (b) The impact on revenue is small when only one collusion group is present, but increases to 50% when 800+ collusion groups are present in an auction system with 4000+ bidders.

2.1 Dynamic Spectrum Auctions

In dynamic spectrum auctions, many small networks compete for a limited range of spectrum. The auctioneer periodically auctions off K spectrum bands to n ($n >> 1$) bidders who submit their bids privately. Without loss of generality, we consider cases where each bidder requests one band and accepts any one of the K bands. Each bidder has a private valuation of the spectrum band derived via either economic modeling [1, 15] or external network survey.

One unique requirement of these new auctions is that the auctioneer must determine auction results based on the interference constraints among bidders. Interfering bidders cannot receive the same spectrum band, while others can reuse the same band. The requirement of "spectrum reuse" leads to two properties of dynamic spectrum auctions:

- **Local competition.** Each bidder competes *locally* with its conflicting neighbors rather than all others. This local competition makes dynamic spectrum auctions fundamentally different from conventional auctions, and imposes significant design challenges.
- **Large scale.** Each spectrum auction can serve a *large* number of bidders with a small set of spectrum bands. Consider an auction with 4000 bidders where each bidder has only up to 5 conflicting peers. In this case, each single spectrum band can allow 1140+ bidders to win the auction without conflicting with each other.

2.2 Bidder Collusion

Collusion occurs in an auction when a group of bidders coordinate their bids to manipulate auction outcomes, gaining unfair advantage. Multiple collusion groups may appear in one auction, but each collusion group is rational – they rig the bids only if this can improve the *group utility*, defined as the sum of individual member's utility [11]. A bidder's *utility* is defined as her true evaluation of the goods minus the price paid if she wins the auction, and otherwise 0.

Despite being legally banned, collusion has appeared in several past commercial auctions, especially those designed without collusion resistance in mind [2, 5, 6, 8, 12]. In those cases, one or multiple collusion groups managed to lower the auction price significantly. Empirical analysis on these auctions, interestingly, reveals that most collusion groups were small (<6 players per group) [2, 6]. This is mainly because they are easy to form but hard to detect in large-scale practical systems. Similarly conclusion was drawn for commercially deployed P2P systems where many collusion groups exist among 160,000+ participants, but the dominant collusion groups were 2–4 players in size [19].

3. COLLUSION IN SPECTRUM AUCTIONS

In this section, we use network experiments to examine the formation and the impact of bidder collusion in emerging dynamic spectrum auctions. We show that the property of local competition provides a fertile breeding ground for collusion, making small-size collusion particularly effective in raising bidder utility and degrading auction revenue.

3.1 Collusion Patterns

We start from identifying representative collusion patterns in large-scale spectrum auctions and examining their effectiveness in raising group utility. Because the pattern depends on the auction design, we use a well-known large-scale spectrum auction design, VERITAS [25], as an illustrative example. This is due to three considerations. First, VERITAS enables spectrum reuse and supports large-scale small networks. Second, VERITAS is a truthful auction that effectively discourages individual cheating. Rational bidders now have no incentive to cheat individually, and will attempt to collude. Third, VERITAS is a representative design of truthful spectrum auctions. It introduces "critical neighbor" based pricing to achieve truthfulness, which has been used widely by recent designs of spectrum auctions [25,26]. In the following, we show that bidders can exploit the local pricing dependency in VERITAS (and many other truthful designs) to form highly effective collusion groups.

We introduce a simple 2-bidder collusion, referred to as the Winner Critical Neighbor (WCN) collusion. Each WCN group contains two members, bidder W and C, where C is the critical neighbor of W in the VERITAS auction. The collusion works as follows: W bids high to maximize its chance to become an auction winner, while C bids extremely low. In VERITAS, an auction winner will be charged by the bid of its critical neighbor. Therefore, if W wins the auction and if C is still his critical neighbor, W will be charged by the bid of C, and the total group utility (W's true valuation minus C's bid) will increase significantly.

WCN collusion group can be formed with relatively low overhead. A bidder only needs to identify its critical neighbor (*e.g.* by studying past auction outcomes). Since interference is a local effect, the critical neighbor must be in its neighborhood. Then the bidder just needs to contact and incentivize the partner via side-payment. In Figure 1 we mark the effective WCN groups on a sample graph that represents the bidder conflict condition. We see that a large number of WCN groups can be formed locally with W and C located in proximity to each other.

To examine the effectiveness of WCN, we simulate a set of VERITAS auctions with 4000 bidders and 1 spectrum band. In each experiment we randomly place the bidders in a given area, and generate the bidder conflict constraints following the graph interference model. With our experimental settings, each bidder sees 5 conflict neighbors in average, mapping to a high degree of spectrum reuse. We assume that without any collusion, the bids are integers randomly distributed in the range of $[1, 100]$. After identifying their critical neighbors, bidders start to form WCN collusion groups where each C reduces her bid to 1 when C is a loser or C's critical value plus 1 (the minimal bid for C to win) when C is a winner. Our experiment results show that in each auction, 800+ WCN groups can effectively increase their group utilities. In Figure 2(a) we plot the group utility improvement of each collusion group when all other groups are colluding in 100 experiments. These results show that many bidders have incentives to collude, since WCN collusion is easy to form and remains highly effective.

3.2 Impact of Collusion

We now examine the impact of collusion from the auctioneer's perspective, focusing on the loss of auction revenue because of collusion. Intuitively, given the scale of the auctions, one would think that small-size collusion will not produce any visible effect. We verify this intuition using the same experiments described in the above. Our main conclusion is that the revenue loss from collusion depends heavily on the number of collusion groups.

- **When one collusion group is present.** Figure 2(b) shows the distribution of revenue loss when a single WCN collusion group is present. We examine the difference in revenue before and after a collusion group rigs the bids. As expected, the loss is almost negligible (<0.3%). This is again because each WCN's impact is local – the original winner now pays significantly less, but the bid change has minimum effect on others.
- **When multiple collusion groups are present.** Figure2(b) also plots the distribution of revenue difference before and after all (800+) WCN collusion groups rig their bids as long as each improves its group utility by colluding. We see that collusion produces significant impact, reducing VERITAS's revenue by 45–50%. In practice, multiple collusion groups are likely coexist, because once bidders notice the incentive to collude, they will independently form colluding groups.

Summary of Findings. Our study shows that the unique requirement of spectrum reuse and resulted local competition provide large incentives for bidders to collude. Even simple, small-size collusion groups of 2-bidders can gain unfair improvements in group utilities. As a result, many small-size collusion groups will form. Although each collusion group leads to small impact, together they will damage the auction revenue and fairness significantly. While our analysis only considered a representative collusion strategy, more collusion patterns will emerge given the complex conflict conditions among bidders, leading to more significant damages. These observations motivate us to find mechanisms that effectively resist bidder collusion, especially small-size collusion.

4. RESISTING COLLUSION

An auction can resist collusion in two ways: either detect and punish colluders harshly, or redesign the auction to *discourage* bidders from colluding. For large-scale spectrum auctions, the detect-and-punish approach is difficult and costly given the complex interference conditions and the network scale. Legally proving collusion did happen is also hard because the collusion is tacit. Therefore, we focus on the second approach.

Proactive Prevention. The concept behind proactive prevention is that (rational) bidders collude only if they can improve the group utility. If an auction is designed to prevent any collusion group from improving their group utility, rational bidders will have little or no incentive to cheat.

Ideally, an auction should diminish the gain of any collusion group, regardless of its *size* and *form of collusive bidding strategy*. Unfortunately, it is proven that the only solution is a trivial posted price method, which can lead to arbitrarily large loss in auction revenue [11]. Therefore, conventional designs have been using a *soft* approach that ensures with a high probability, no collusion group can gain in group utility [11, 20]. The probability, in general, depends on the size of the collusion group and the auction design.

Inspired by [11, 20], we consider the soft approach in designing collusion-resistant spectrum auctions. Our design, however, faces two new challenges. First, existing designs [11, 20] do not consider heterogeneous bidder conflict constraints and local competition. When being directly applied or extended to spectrum auctions, they either suffer severe interference or lose the desired collusion resistance. Second, existing designs [11, 20, 23, 24] often require the optimal allocation to enforce collusion resistance. Distributing spectrum with spatial reuse, however, is a NP-hard problem. In large-scale spectrum auctions, we can only consider approximate solutions. Unfortunately, it has been shown that many prior works lose their desired economic property by replacing the optimal allocation with an approximate solution [4, 13, 23–25].

A Case for Divide and Conquer. A closer look at the problem shows that the above challenges are triggered by the complex bidder conflict constraints. To overcome these challenges, we propose to *decouple* the collusion-resistance design from the complex interference constraints, using the concept of *divide and conquer*. First, we divide bidders into different segments such that bidders in each segment are free of interference constraints and can use a single spectrum band simultaneously. Second, inside each segment, the auction process (*i.e.* determining prices and winners) falls back to that of the conventional ones and can be "conquered" using existing solutions [11, 20]. Finally, the auctioneer selects K winning segments, one for each band, and assigns winners in each segment a spectrum band. Via the decoupling, this design successfully reduces a complex auction design into a form that can be handled by conventional solution.

This design, however, cannot fully address collusion, because it only enforces collusion resistance within each segment but not globally. It is effective towards any collusion group involving bidders from the same segment (intra-segment collusion), but cannot diminish gain of a collusion group involving bidders from different segments (inter-segment collusion). Therefore, we need an extra layer of collusion resistance to effectively suppress both intra- and inter-segment collusion. This is the basic idea behind Athena, which we will introduce next in detail.

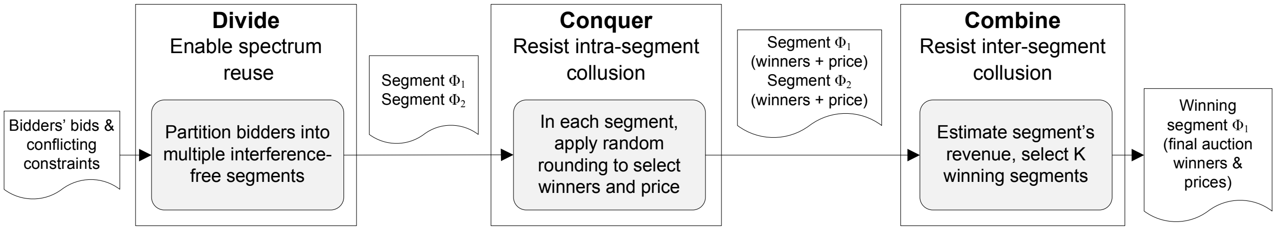

Figure 3: Athena's 3-stage decision process: Divide, Conquer, and Combine.

5. ATHENA

We propose Athena, a new collusion resistant spectrum auction design. Like [11, 20], Athena applies soft collusion resistance to address small-size collusion, ensuring that with a very high probability, no small-size collusion groups gain from cheating. Different from [11, 20], Athena operates under the complex bidder interference constraints and exploits spectrum reusability to serve a large number of bidders. Different from [23, 24], Athena resists *any* collusive bidding strategies using computationally efficient algorithms, making it a low-cost and deployable solution.

Athena uses a 3-stage process: Divide, Conquer, and Combine (Figure 3). It first applies a spectrum allocation algorithm in Divide to enable spectrum reuse, and then applies two layers of collusion-resistant designs in both Conquer and Combine to resist intra-segment and inter-segment collusion. Next we describe Athena's design towards its two goals: *enabling spectrum reuse* and *diminishing collusion gain*. We also define and prove Athena's collusion resistance.

5.1 Enabling Spectrum Reuse

After collecting bids, Athena first divides bidders into multiple non-overlapping segments based on their interference constraints. Bidders in each segment do not conflict with each other and can use the same spectrum band simultaneously without any interference. To do so, Athena uses a spectrum allocation algorithm to *virtually* assign one spectrum band to bidders without limiting the number of bands used. It then groups bidders assigned to the same band into a segment. An important requirement is that the bidder partition must be independent of their bids. Otherwise bidders can rig bids to manipulate the segment formation and the auction result.

The outcome of this stage is a set of V bidder segments, $\Phi_1, \Phi_2, ..., \Phi_V$, where V depends only on the underlying bidder conflict constraints, but not their bids or the number of spectrum bands to be auctioned.

5.2 Diminishing Collusion Gain

The pattern of collusion can be divided into two categories: intra-segment collusion formed by members in the same segment, and inter-segment collusion formed by those in multiple segments. To resist both types, Athena embeds a collusion-resistance design into both Conquer and Combine, where Conquer determines potential winners and prices in each segment, and Combine selects winning segments.

Addressing Intra-segment Collusion. In the first step, Athena processes each segment individually and selects potential winners, while diminishing the impact of any local collusion group. For each segment, Athena determines a price, and selects bidders bidding no less than this price as the potential winners. Because bidders in each segment are free of interference constraints, enforcing collusion-resistance in each segment is much simper and can use existing solutions [11, 20]. Focusing on soft resistance, we use the t-Truthful with Probability (tCP) solution from [11]. The general idea of tCP is to randomize the winner selection such that the result becomes relatively insensitive to bid changes. In this case, small-size collusion groups (of size t or less), with a high probability p, cannot rig their bids to change the pricing. For each segment Φ_m, its p is a function of t and the selection configuration (c_m, α_m). We briefly outline the procedure of tCP in the Appendix. The detailed procedure can be found in [11].

The outcome of this step is that each segment Φ_m produces a price Γ_m and a set of $N(\Gamma_m)$ potential winners who bid no less than Γ_m. These winners will become auction winners and be charged with Γ_m if the segment is selected in the next step.

Addressing Inter-segment Collusion. In this step, Athena selects winning segments while diminishing the impact of any collusion group built across segments. Because only K out of V $(K \leq V)$ segments will be assigned with a spectrum band, Athena must carefully select winning segments so that no collusion group can rig their bids to manipulate the auction outcome.

Intuitively, the selection should be straightforward. To maximize auction revenue, Athena should pick segments that produce the highest revenue. That is, Athena should rank each segment Φ_m by its potential revenue $\Gamma_m \cdot N(\Gamma_m)$, and select the K segments with the largest revenue. This design, however, cannot fully address collusion. We prove this using a sample collusion pattern of two bidders.

Consider bidders A$\in \Phi_1$ and B$\in \Phi_2$. When they do not collude and submit their original bids, Φ_1 uses a price $\Gamma_1 = 10$ and has 201 winners, and Φ_2 uses a price $\Gamma_2 = 5$ and has 401 winners. Thus Φ_1 is the winner by having higher revenue $(10 \cdot 201 \geq 5 \cdot 401)$. Now assume A reduces its bid from 11 to 9. Although it does not change Φ_1's price Γ_1, the number of winners in Φ_1 reduces to 200 and Φ_2 becomes the winner. Now B becomes the winner and increases its utility from 0 to 15. B can easily transfer side-payment to incentivize A to set up the collusion.

The above collusion is effective because bidders can manipulate the potential revenue produced by each segment and hence affect the segment selection result for their own gains. As long as bidders know that they can gain in group utility by colluding, they will be incentivized to collude even though they may not know the optimal collusive strategy in most cases. However, such random collusion from multiple groups can damage the auction revenue significantly.

Athena overcomes this vulnerability using a new method to estimate the segment revenue. Specifically, Athena replaces the original absolute revenue $\Gamma_m \cdot N(\Gamma_m)$ with an estimated revenue:

$$\hat{R}(\Gamma_m) = \Gamma_m \times g_c(N(\Gamma_m)), \tag{1}$$

where $g_c(\cdot)$ is a randomized rounding function parameterized by c to randomly round $N(\Gamma_m)$. The rounding ensures that $\forall x > 0$ and $\forall y \in [x-t, x+t]$, with probability $(1-\log_c \frac{x+t}{x-t})$, $g_c(y)$ only depends on x and t, but not y [10]. That is, the rounding makes $\hat{R}(\Gamma_m)$ insensitive to both $N(\Gamma_m)$ and Γ_m that can be affected by collusive bids. Because the estimated revenue is only used to select winning segments, the impact of such rounding is minimum.

5.3 Athena's Collusion Resistance

Athena achieves the following soft collusion resistance:

Definition 1. *An auction achieves the (t,p)- **truthfulness** if with a probability of p or higher, no collusion group of size t or less can improve its group utility by rigging the bids. This holds even if multiple collusion groups are present, as long as each group is of size t or less.*

Theorem 1. *Athena achieves the (t,p)-truthfulness with $p = 1+\log_{c_{min}}(1-\lambda t/(l^{min}-t))$ where l^{min} is the number of winners of the smallest segment that runs tCP, c_{min} and λ are auction parameters. When $t/l^{min} << 1$, $p = 1 - O(t/l^{min})$.*

The proof is in the Appendix. In the above, $\lambda = \frac{2c^{max}\alpha^{min}}{\alpha^{min}-1}$, $c^{min} = \min\{c_1, ..., c_V\}$, $c^{max} = \max\{c_1, ..., c_V\}$, $\alpha^{min} = \min\{\alpha_1, ..., \alpha_V\}$ when all V segments run tCP. And (c_m, α_m) are the tCP configuration in segment Φ_m.

One restriction in Athena is that any segment Φ_m, in order to run tCP, must be large enough to ensure $\alpha_m > 1$ [11]. After some computation, we can formally define the restriction as follows: if all V segments run tCP, then the segment with the least number (l^{min}) of winners must have more than l_{tCP} winners, where

$$l_{tCP} \triangleq \frac{2c^{max}t}{1-(c^{min})^{p-1}} + t > t. \qquad (2)$$

Hence we have $l^{min} > l_{tCP} > t$. In some cases, certain segments cannot satisfy the l_{tcp} requirement. For these segments, Athena uses posted price to select winning bidders. Because posted price achieves hard collusion-resistance, it is easy to show that Athena's collusion resistance only depends on segments that run tCP.

6. FINE-TUNING ATHENA

Having demonstrated Athena's collusion resistance, in this section we focus on fine-tuning Athena to improve its auction revenue. The configuration procedure is independent of bids so that bidders cannot manipulate their bids to affect the configuration results and hence the auction outcomes.

6.1 Athena with Uniform Segments

We start from the simplest case where segments are of the same size, so they either all run tCP or all run posted price to decide winners in each segment. Intuitively, one would prefer tCP to posted price if the l_{tCP} constraint is met, because tCP guarantees a revenue with a distance of $c_m\alpha_m$ to the optimal [10]. When $c_m\alpha_m$ is high, however, using tCP could produce lower *average* revenue than that of posted price. When an auction uses the average revenue as its performance criterion, it should carefully configure c_m,α_m in each segment m running tCP. In this simple case, because the segments are of equal length, and that c_m and α_m do not depend on the bids, we have $c_m = c^{min} = c^{max}$, and $\alpha_m = \alpha^{min}$, $\forall m$. From [10], the value of c_m depends only on the size of the segment and the resistance level t. From Theorem 1, the choice of α_m depends on t, p, c^{min}, c^{max} and l^{min}.

Algorithm 1 Athena-Configuration(t, p, $\mathbb{N}$, Y)

Input: 1) (t,p) requirement; 2) $Y = \min(V, K)$ for V segments and K channels; 3) Y largest segments sizes $\mathbb{N} = \{N_i | i \leq Y\}$.
Output: (l_{tCP}, α_{tCP})

Rank segments by sizes: $N_1 \geq N_2 ... \geq N_Y$
for $m = 1$ to Y **do**
  Set $l^{min} \leftarrow l_m = \lceil \frac{N_m}{2} \rceil$, $c^{max} \leftarrow c_m$ by (3), $c^{min} \leftarrow c_1$
  Set $\alpha_m = \alpha^{min}$ by (4), l_{tCP} by (2)
  if $(l_m < l_{tCP})$ **or** $(\alpha^{min} \geq \theta_{tCP})$ **then**
    //Estimate revenue when all run posted price
    $E(m) = \sum_{i=1}^{Y} E^P(N_i), \quad c_m = 0$
  else
    //Estimate revenue when m segments run tCP
    $E(m) = \sum_{i=1}^{m} E^{tCP}_{\alpha^{min}}(N_i) + \sum_{i=m+1}^{Y} E^P(N_i)$
  end if
end for
$m^* = \text{argmax}_m E(m)$
if $c_{m^*} > 0$ **then**
  $(l_{tCP}, \alpha_{tCP}) = (l_{m^*}, a_{m^*})$
else
  All run posted price
end if

Given such dependency, we configure Athena as follows:

1. By [10], configure c_m as
$$c_m = \text{argmax}_x[(\frac{l_m - t}{l_m + t} - \frac{1}{x})/\ln(x)], \qquad (3)$$
where the number of winners l_m is estimated as half number of bidders in the segment. Compute l_{tCP} as (2), if $l_m \geq l_{tCP}$, go to step 2, otherwise run posted price.

2. By Theorem 1, use c_m to derive α^{min}; set $\alpha_m = \alpha^{min}$.
$$\alpha^{min} = \frac{(1-(c^{min})^{p-1})(l^{min}-t)}{(1-(c^{min})^{p-1})(l^{min}-t) - 2c^{max}t} \triangleq \alpha_{tCP}. \qquad (4)$$

3. Compute a threshold θ_{tCP} where if $\alpha_m \geq \theta_{tCP}$, then using tCP with (c_m, α_m) will *in average* lead to lower revenue than that using posted price. Given the size of the segment, the value of θ_{tCP} is derived statistically using a bid distribution. To ensure collusion resistance, the choice of θ_{tCP} cannot depend on the actual bids. In addition, since the auctioneer in general does not have good knowledge of the bid distribution, Athena assumes the uniform random distribution and computes the expected revenue of a segment with size x as $E^{tCP}_\alpha(x) = \alpha^{-1}(1-\alpha^{-1})x$ when it runs tCP using α, and $E^P(x) = x/6$ when it runs posted price. From these we can derive $\theta_{tCP} = (3+\sqrt{3})$. The detailed derivations are omitted due to the space limitations.

4. If $\alpha_m < \theta_{tCP}$, then all the segments use tCP with (c_m, α_m), otherwise they use posted price.

The configuration is independent of the actual bids. This is necessary to ensure collusion resistance, preventing bidders from manipulate the mechanism used in each segment thus the price and winner selection.

6.2 Athena with Non-Uniform Segments

When segments are of different sizes, the configuration becomes more complex due to the interdependency across segments. For all the segments that run tCP, we will set $\alpha_m = \alpha_{tCP}$ to maximize the auction revenue. But this also means that the value of α_{tCP} is constrained by the smallest segment running tCP. By including a smaller segment, l^{min} decreases and α_{tCP} increases (according to (4)), thus the revenue from these segments decreases. Hence we need to judiciously set the virtual clearing scheme in each segment.

The principle guiding Athena configuration is the fact that the choice of segments running tCP is monotonic in size. If a segment of size L runs tCP, then a segment of size larger than L should also run tCP. Using this principle, Athena first sorts the segments by their sizes, and searches for the best number of segments running tCP that produces the best average revenue. For K segments, there are only K choices, so that the computation is of low complexity. Algorithm 1 lists the detailed process. Again the configuration is independent of the bids to ensure collusion resistance.

The interdependency across segments also creates a dilemma when forming segments. Creating balanced segments will allow more segments to use tCP but also suffer from a larger α_m and lower revenue guarantee. On the other hand, imbalanced partition will have less segments running tCP but those that run tCP are large in size and thus benefit from a lower α_m and higher revenue performance. In Section 7.3, we will study the impact of segment partition in more detail.

6.3 Athena's Revenue Bound

For a given segment partition and assuming all the segments are large enough to run tCP, we show that Athena's revenue is with a bounded distance to the optimal. The revenue bound is stated as the following theorem. We include the proof in the appendix.

Theorem 2. *Given a segment partition, let S denote the set of segments running tCP in Athena. For all segments in S, while satisfying the (t,p)-truthfulness, Athena achieves an auction revenue no less than $R^{OPT}/((c^{max})^2\alpha_{tCP})$, where R^{OPT} is the sum of the optimal revenue obtained by treating each segment $\Phi_m \in S$ separately, c^{max} and α_{tCP} are the auction parameters required to achieve the (t,p)-truthfulness, as defined in Theorem 1 and (4).*

The above theorem indicates that c^{max} and α_{tCP} are the two dominating factors in Athena's revenue. From Section 6.1, their values primarily depend on the minimum segment size l^{min} and the choice of (t,p). Take the case of uniform segments as an example, Table 1 lists the values of c^{max} and α_{tCP} under various network configurations. We see that both c^{max} and α_{tCP} decrease as the segment size increases. This is expected because Athena uses a statistical method to control collusion gain. The larger the segment size, the better the efficiency. In the case of non-uniform segments, the values of c^{max} and α_{tCP} are similar to that of uniform segments for each given l^{min}.

l^{min}	$(t,p)=(2,0.8)$		$(t,p)=(2,0.9)$		$(t,p)=(4,0.8)$	
	c^{max}	α_{tCP}	c^{max}	α_{tCP}	c^{max}	α_{tCP}
540	1.1998	2	1.1998	3.3732	1.3636	3.1096
600	1.188	2	1.188	3.1503	1.2828	2.4792
800	1.1597	2	1.1597	2.6965	1.2381	2.2059
1000	1.141	2	1.141	2.4463	1.2089	2.0478

Table 1: c^{max} and α_{tCP} under various configurations.

6.4 Athena's Computational Complexity

Besides the spectrum allocation algorithm used in the Divide step, Athena requires only polynomial time to run the rest of its components. The following theorem summarizes Athena's computational complexity.

Theorem 3. *Given n bidders that form V segments and compete for K spectrum bands, the overall complexity of Athena is $O(n\log(n))$ on top of the complexity of the spectrum allocation algorithm used to partition bidders.*

Proof. Athena's complexity comes from the 3-stage procedure and the fine-tuning (Algorithm 1). Assume that n bidders form V segments and bid for $K(\leq V)$ channels. First, in the main procedure of Athena, the complexity of Divide depends on the spectrum allocation algorithm; the complexity of Conquer is from rounding the bids hence is linear to the number of the bids $O(n)$; and the Combine stage takes $O(V\log(V))$ time to sort the segments. Second, in the fine-tuning of Athena, it takes $O(V\log(V))$ to sort the segments' sizes and $O(V)$ to find the statistically optimal parameter configuration. Thus by $V \leq n$, the overall complexity of Athena is $O(n\log(n))$, plus the complexity of the spectrum allocation algorithm. Most existing allocation algorithms [21, 22] are of polynomial-time complexity and can be used in Divide. □

7. ATHENA EXPERIMENTS

Having verified Athena's collusion resistance and bounded revenue analytically under any bids, in this section we use two case studies to examine both properties experimentally. In the first study, we verify Athena's collusion resistance by identifying some collusion patterns that are effective in a truthful auction (*e.g.* VERITAS [25]), and examining their effectiveness in Athena. In the second study, we examine the cost of collusion resistance as the revenue loss in "ideal" cases. When bidders are aware of Athena's robustness, they have no incentive to cheat but bid their true valuations. We define the cost of collusion resistance as the difference between Athena's revenue and that of VERITAS [25] (without collusion resistance). We study the cost experimentally using valuation distributions verified by past auction field experiments [14].

Our experiments consider large spectrum auctions with 4000 randomly deployed bidders; each bidder in average sees 5 conflicting peers. For a fair comparison between Athena and VERITAS, we use the greedy spectrum allocation algorithm of VERITAS to divide bidders into segments. Because VERITAS's revenue depends heavily on the number of bands auctioned K [25], we configure K to the one that produces the highest revenue for VERITAS. In many cases, $K=2$ and supports 2200+ winners. In the second case study, we also compare Athena with posted price [11], which achieves hard collusion resistance at significant revenue loss.

When configuring Athena, we consider $p \in [0.7, 1]$ and t=2, 4, 8 to focus on collusion of small group size. Note that t is the largest size of collusion groups that the system seeks to resist, but not the total number of colluders.

7.1 Collusion Resistance

We consider the WCN collusion that is shown to be highly effective in VERITAS (see Section 3). For 100 auction trials each with different bidder interference constraints, we first identify effective WCN collusion groups in VERITAS. We then deploy them in Athena and examine the group utility improvement after they rig the bids.

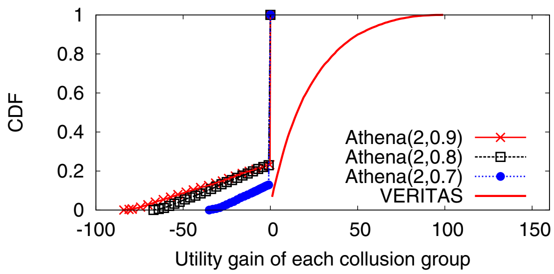

Figure 4: Athena resists collusion by preventing the collusion group from improving the group utility.

Figure 4 plots the cumulative distribution of the group utility improvement for both Athena and VERITAS. We make two observations. First, different from VERITAS, in Athena the WCN collusion leads to no gain but significant loss in group utility. This is because W and C in each WCN group are located in different segments. Since only C lowers its bid, W's utility remains the same yet C's utility most likely becomes zero by bidding very low, and the group utility drops significantly. Second, the loss increases with the level of collusion resistance, p in this case. This is because the random rounding in Athena becomes coarser as p increases, leading to a lower price in each segment and higher utilities for winners. Hence the group utility loss (due to C not winning) varies in larger ranges. Combined with the analytical proof (Theorem 1), this experimental result demonstrates the effectiveness of Athena's collusion resistance.

7.2 Cost of Collusion Resistance

With Athena's collusion resistance, bidders have little incentive to collude, but will likely bid by their true valuation of the spectrum. Under this stable condition, we examine the cost of such collusion resistance by comparing the revenue of Athena to that of the idealized VERITAS assuming no one cheats. Athena's loss in revenue in this ideal case represents the necessary cost to ensure collusion resistance.

In Figure 5(a)-(b) we examine the cost in terms of the cumulative distribution of Athena's normalized revenue loss over VERITAS, and in Figure 5(c) we examine the average loss of revenue over different (t, p). The results are derived over 5000 bid (valuation) generations assuming the bids follow the uniform random distribution over [1,100]. As a reference, we include posted price's performance as well as the revenue loss in VERITAS due to the WCN collusion.

Results in Figure 5 lead to two key observations. First, compared with posted price, Athena's collusion resistance comes at a much lower and "stabler" cost. The large variance in posted price's revenue cost also confirms the analytical conclusion where it does not provide any revenue guarantee but leads to arbitrarily large revenue loss. Second, Athena's cost for collusion resistance scales gracefully with the level of resistance (t, p). Compared to VERITAS, Athena sacrifices 5%, 14% and 21% of the revenue to achieve collusion-resistance with $t \leq 2, 4, 8$, and $p = 0.8$, and the cost increases to 18%, 28% and 33% when $p = 0.9$. Hence Athena is more sensitive to the value of p. Nevertheless, the revenue reduction is still much smaller than VERITAS's loss due to collusion (50+%). This example demonstrates Athena's cost-effectiveness, but also implies that the auctioneer needs to carefully choose the (t, p) configuration. We defer this question to a future study.

We also repeat the study using non-uniformly distributed bids (valuations). We use the beta distribution, one of the most popular models on bidder behaviors [14]. It represents scenarios where the valuation is randomly distributed around a commonly known value with bounded support. Results in Figure 6 show a similar trend between Athena and posted price. But Athena's cost in revenue reduces significantly compared to the uniform distribution. This is because Athena and VERITAS react differently to bid distribution. Athena applies a single-price scheme based on random rounding in each segment, creating more winners when the bids are clustered. In contrast, VERITAS's revenue is driven by the competition level, and clustered bids will not increase the number of winners.

Gain from Athena's Fine Tuning. We also examine how Athena's fine-tuning configuration helps reduce the cost of collusion resistance. We compare it with the basic configuration that applies tCP in all the qualified segments without optimizing α globally. Figure 7 plots the ratio of Athena's revenue over that of the basic configuration under uniform and non-uniform bid distributions with $(t,p) = (2, 0.9)$. Because the basic configuration allows all qualified segments to run tCP including those small ones, the revenue is constrained by the smallest segment. By judiciously choosing the segments to run tCP, Athena's fine-tuning increases revenue by 25% in average under both bid distributions.

7.3 Impact of Segment Formation

Our analytical results show that Athena's performance depends heavily on the segment sizes. The sizes depend on both the spectrum allocation algorithm and the bidder interference constraints, which we will exploit next.

We first consider three representative spectrum allocation algorithms: Max-IS [22], Greedy [21] and random allocation. In principle, RAND $\leq$ Greedy $\leq$ Max-IS in terms of the allocation efficiency. Yet when combined with Athena, their difference becomes minimum (<5%). This is because all three algorithms produce similar-sized large segments, and differ mostly in their small, medium-sized segments. Since Athena's performance depends heavily on the large segments that run tCP, the difference among these algorithms is small.

Next, we examine whether it is always beneficial to have balanced segments, by assuming that bidders can be partitioned arbitrarily and we can produce arbitrary segment patterns. Figure 8 compares three randomly generated partitions of size [3500, 400, 50, 20, 10, 10, 5], [2000, 900, 500, 300, 200, 50, 50], and [1500, 800, 500, 300, 200, 200, 150, 150, 100, 50, 50], with increasing balance across segments. We also plot the ideal case where all 4000 bidders are in one segment. Aside from the ideal case, Athena favors the most imbalanced partition, particularly when the spectrum is limited. This is because balanced partition often leads to smaller l^{min}, thus larger c^{max} and α_{tCP}. For example, when the number of channels auctioned $K = 2$, (c^{max}, α_{tCP}) is $(1.0965, 2)$ in the ideal case, but it increases to $(1.2381, 2.8006)$ under the [1500, 800,...] partition.

The above results do not tell us in absolute terms what the best spectrum allocation algorithm is for Athena. Relatively speaking, however, we see that creating larger unbalanced segments provides better revenue. We leave the quest for the best algorithm to a future study. A nice property of Athena is that it can use *any* spectrum allocation algorithm as long as it produces segments with conflict-free bidders.

8. RELATED WORK

Collusion resistance has been widely studied in conventional auctions. Prior work proposes solutions like tCP [11,

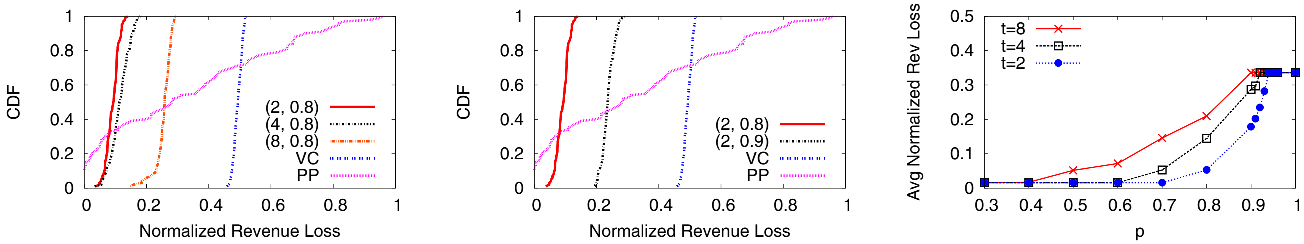

(a) Varying t, uniform bid distribution. (b) Varying p, uniform bid distribution. (c) Average normalized revenue loss.

Figure 5: Cost of Athena's collusion resistance assuming uniform bid distribution. (a)-(b):CDF of the normalized revenue loss over VERITAS, where PP denotes posted price, and VC denotes VERITAS under the WCN collusion. (c):the average normalized revenue loss over VERITAS.

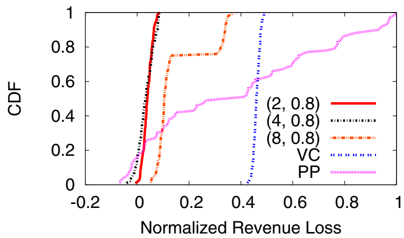

Figure 6: Cost of Athena's collusion resistance, non-uniform bid distribution.

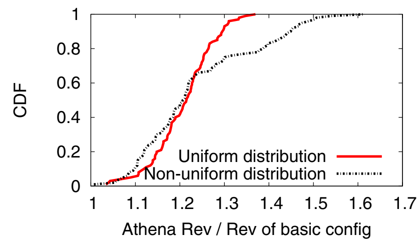

Figure 7: The gain of Athena's fine tuning over the basic configuration.

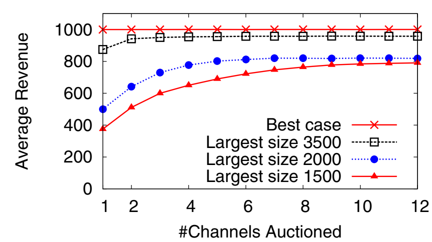

Figure 8: Impact of segment formation, $(t, p) = (4, 0.8)$.

20] for auctions where bidders either all conflict or do not conflict at all. Athena utilizes tCP to resist intra-segment collusion, but focuses on designing a general framework for dynamic spectrum auctions with general bidder interference constraints. Recent work [16] proposes a pricing game for spectrum auctions but does not consider any spectrum reuse.

Another work considers collusion-resistance with spectrum reuse [23, 24]. It, however, targets small-scale auctions and only considers three specific types of collusion strategy. In practice, collusion behaviors are highly complex and hard to predict, especially in large-scale auctions. Different from [23, 24], Athena provides a general collusion resistance that can address any form of collusive bidding strategy. Perhaps more importantly, [23, 24] requires solving a NP-hard optimization to ensure its partial collusion resistance. Athena, on the other hand, can use any approximate spectrum allocation solution (of polynomial-time complexity), and yet still maintains its general collusion resistance.

9. CONCLUSION AND FUTURE WORK

We propose Athena, a new spectrum auction design to resist small-size bidder collusion. Using the concept of "Divide and Conquer," Athena decouples the problem of spectrum allocation from that of economic mechanism design, achieving spectrum reuse, revenue guarantee, and soft collusion-resistance. To our best knowledge, Athena is the first to address any form of small-size collusive bidding in large-scale spectrum auctions with spectrum reuse.

There are several directions to extend Athena. First, Athena assumes that each bidder requests one channel. It is worthwhile to extend Athena to allow bidders to request multiple channels. Multi-channel bidding, however, complicates collusion strategy, expanding its impact on auction outcomes. One must carefully analyze the collusion behavior and take additional precaution to suppress collusion gain. Second, Athena can use any spectrum allocation algorithm to form segments, yet it is beneficial to identify good allocation algorithms that lead to higher auction revenue.

Acknowledgments

We thank our shepherd Ranveer Chandra, and the anonymous reviewers for their helpful suggestions. This work is supported in part by NSF Grants CNS-0915699 and CNS-0832090. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

10. REFERENCES

[1] Bajari, P., and Fox, J. T. Complementarities and collusion in an FCC spectrum auction. Tech. Rep. 11671, National Bureau of Economic Research, October 2005.

[2] Bajari, P., and Yeo, J. Auction design and tacit collusion in FCC spectrum auctions. Tech. Rep. 14441, National Bureau of Economic Research, October 2008.

[3] Buddhikot, M., and Ryan, K. Spectrum management in coordinated dynamic spectrum access based cellular networks. In *Proc. of DySPAN* (2005).

[4] Clarke, E. Multipart pricing of public goods. *Public Choice XI* (1971), 17–33.

[5] Cramton, P., and Schwartz, J. Collusive bidding: Lessons from the FCC spectrum auctions. *Journal of REgulatory Economics 17* (May 2000), 229–252.

[6] Cramton, P., and Schwartz, J. Collusive bidding in the FCC spectrum auctions. Tech. rep., U. Maryland, 2002.

[7] Fleishman, G. So many nodes, no little security. *The New York Times* (2002).

[8] Friedman, M. Comment on 'collusion in the auction market for treasury bills'. *J. of Political Economy 9* (1996), 757–785.

[9] Gandhi, S., et al. A general framework for wireless spectrum auctions. In *Proc. of DySPAN* (2007).

[10] Goldberg, A. V., and Hartline, J. D. Competitiveness via consensus. In *Proc. of SODA* (2003).

[11] Goldberg, A. V., and Hartline, J. D. Collusion-resistant mechanisms for single-parameter agents. In *Proc. of SODA* (2005).

[12] Goswami, G., Noe, T. H., and Rebello, M. J. Collusion in uniform-price auctions: Experimental evidence and implications for treasury auctions. *Review of Financial Studies 72* (1964), 513–514.

[13] Groves, T. Incentive in terms. *Econometrica 41* (1973), 617–631.
[14] Gupta, A., and Nadarajah, S. *Handbook of Beta Distribution and Its Applications.* CRC Press, 2004.
[15] Ileri, O., and Mandayam, N. Dynamic spectrum access models: toward an engineering perspective in the spectrum debate. *IEEE Commu. Magazine 46*, 1 (2008), 153–160.
[16] Ji, Z., and Liu, K. Multi-stage pricing game for collusion-resistant dynamic spectrum allocation. *IEEE JSAC 26*, 1 (Jan. 2008), 182–191.
[17] Jia, J., Zhang, Q., Zhang, Q., and Liu, M. Revenue generation for truthful spectrum auction in dynamic spectrum access. In *Proc. of MobiHoc* (2009).
[18] Krishna, V. *Auction Theory.* Academic Press, March 2002.
[19] Lian, Q., et al. An empirical study of collusion behavior in the maze P2P file-sharing system. In *Proc. of ICDCS* (2007).
[20] McSherry, F., and Talwar, K. Mechanism design via differential privacy. In *Proc. of FOCS* (2007).
[21] Ramanathan, S. A unified framework and algorithm for channel assignment in wireless networks. *Wirel. Netw. 5*, 2 (1999), 81–94.
[22] Subramanian, A. P., et al. Fast spectrum allocation in coordinated dynamic spectrum access based cellular networks. In *Proc. of DySPAN* (2007).
[23] Wu, Y., et al. Collusion-resistant multi-winner spectrum auction for cognitive radio networks. In *Proc. of Globecom* (2008).
[24] Wu, Y., Wang, B., Liu, R., and Clancy, T. A scalable collusion-resistant multi-winner cognitive spectrum auction game. *IEEE Transactions on Communications* (Dec. 2009).
[25] Zhou, X., et al. eBay in the sky: Strategy-proof wireless spectrum auctions. In *Proc. of MobiCom* (2008).
[26] Zhou, X., and Zheng, H. TRUST: A general framework for truthful double spectrum auctions. In *Proc. of INFOCOM* (2009).

Appendix

Preliminary of tCP. tCP applies to auctions where bidders do not conflict with each other. Let $\mathbb{B}_m$ represent the bid set. With tCP, the auctioneer performs the following procedure to choose a price Γ_m and sets the bidders who bid no less than Γ_m as winners. It starts from choosing a parameter $\alpha > 1$ and defining $\mathbb{G} = \{\alpha^i | i \in \mathbb{Z}\}$ as the set of candidate prices. For each price candidate $\alpha^i \in \mathbb{G}$, let $N(\alpha^i)$ be the number of bids no less than α^i in $\mathbb{B}_m$. Next, tCP introduces a *consensus estimation* function $g_c(\cdot)$ parameterized by c to randomly round $N(\alpha^i)$. It then computes the price $\Gamma_m = \arg\max_{\alpha^i \in \mathbb{G}} \alpha^i \cdot g_c(N(\alpha^i))$. The random rounding $g_c(\cdot)$ makes the choice of price insensitive to bid changes, ensuring that with high probability, Γ_m will not be affected by no more than t bids. As proved in [11], tCP achieves the (t,p)-truthfulness, where p is lower bounded by Lemma 1.

Lemma 1. *A tCP auction with parameters (c, α) is (t,p)-truthful with $p = 1 + \log_c(1 - \frac{\lambda t}{l-t})$, where $\lambda = \frac{2c\alpha}{\alpha - 1}$, and l is the number of winners when there is no collusion.*

Proof of Theorem 1. We consider the scenario where the auction contains multiple segments. Lemma 1 directly applies when there is a single segment.

Proof. In "Divide," bidders cannot rig bids to change the segment formation because the allocation is bid-independent. This ensures the first level of collusion-resistance.

Next in "Conquer", assume there are V segments where each Φ_m uses tCP with (c_m, α_m) for virtual clearing and has $_m$ winners $(m = 1, ..., V)$. $c^{min}, c^{max}, \alpha^{min}, l^{min}$ are defined as Section 5.3. Consider any collusion group of size $\leq t$. Let t_m be the number of its members assigned to segment Φ_m, $\sum_{m=1}^{V} t_m \leq t$. We introduce the notion of *t-truthful*: an auction is *t-truthful* if it can diminish the gain of all forms of collusion groups of size t or less. Thus an (t,p)-truthful auction means that it is t-truthful with a probability of p or higher. Let Pr_m^{tT} be the probability that Φ_m is t_m-truthful. By Lemma 1 with $\lambda_m = \frac{2c_m\alpha_m}{\alpha_m - 1}$:

$$Pr_m^{tT} \geq 1 + \log_{c_m}(1 - \frac{\lambda_m t_m}{l_m - t_m}), \tag{5}$$

$$\widetilde{Pr}_m^{tT} = 1 - Pr_m^{tT} \leq -\log_{c_m}(1 - \lambda_m \frac{t_m}{l_m - t_m}). \tag{6}$$

By Athena's "Combine", if each segment Φ_m is t_m-truthful, the overall auction is t-truthful because collusive bids cannot affect any Φ_m's bid $\hat{R}(\Gamma_m)$ and thus Athena's auction result. Let Pr^{tT} be the probability that Athena is $t-$truthful, we have

$$\begin{aligned} Pr^{tT} &\geq \prod_{m=1}^{V} Pr_m^{tT} = \prod_{m=1}^{V}(1 - \widetilde{Pr}_m^{tT}) \geq 1 - \sum_{m=1}^{V} \widetilde{Pr}_m^{tT} \\ \text{(from (6))} &\geq 1 + \sum_{m=1}^{V} \log_{c_m}(1 - \lambda_m \frac{t_m}{l_m - t_m}). \end{aligned} \tag{7}$$

Since $\log_{c_m}(1 - \lambda_m \frac{t_m}{l_m - t_m}) < 0$, its value decreases as c_m decreases. We have:

$$\begin{aligned} Pr^{tT} &\geq 1 + \sum_{m=1}^{V} \log_{c^{min}}(1 - \lambda_m \frac{t_m}{l_m - t_m}) \\ &\geq 1 + \log_{c^{min}}(1 - \sum_{m=1}^{V} \lambda_m \frac{t_m}{l_m - t_m}). \end{aligned} \tag{8}$$

Let $\lambda = \frac{2c^{max}\alpha^{min}}{\alpha^{min} - 1}$. By the definitions of c^{max} and α^{min}, we have $\lambda_m = \frac{2c_m\alpha_m}{\alpha_m - 1} \leq \lambda$. Using the property of function $\log(1 - x\frac{y}{z-y})$, we reduce (8) into

$$\begin{aligned} Pr^{tT} &\geq 1 + \log_{c^{min}}(1 - \lambda \frac{\sum_{m=1}^{V} t_m}{l^{min} - t}) \\ &\geq 1 + \log_{c^{min}}(1 - \lambda \frac{t}{l^{min} - t}). \end{aligned} \tag{9}$$

By Definition 1, we see that the auction is (t,p)-truthful with $p = 1 + \log_{c^{min}}(1 - \lambda t/(l^{min} - t))$. When t/l^{min} is very small, this bound is approximately $p = 1 - \mathrm{O}(t/l^{min})$. □

Proof of Theorem 2. We prove Athena achieves bounded revenue when tCP is applied in all segments.

Proof. For each segment Φ_m, let $\hat{R}_m$ be the estimated revenue as (1) in Athena, R_m be the revenue achieved by Athena if Φ_m is chosen, and R_m^{OPT} be the optimal revenue achieved by OPT in Φ_m. By [11], we have

$$R_m \geq \hat{R}_m \geq R_m / c_m, \tag{10}$$

where c_m is defined in Section 5.3. Given K channels, only $\min(K, V)$ segments can be final winning segments. Since Athena chooses winning segments based on their estimated revenue (See (1)), Athena and OPT may choose different segments. Let S^{OPT}, S^{Athena} be the set of segments chosen by OPT and Athena respectively, we have $|S^{OPT}| = |S^{Athena}|$. We can then compute c^{max} and α_{tCP} based on (t,p) and S^{Athena}.

We divide S^{Athena} into two sets Ψ_1 and Ψ_2, where $\Psi_1 = S^{OPT} \bigcap S^{Athena}$, and $\Psi_2 = S^{Athena} \setminus \Psi_1$. Next we prove the revenue of Athena in both subsets is bounded.

First, for $\Phi_m \in \Psi_1$, by [11], we have $R_m \geq R_m^{OPT}/(c_m\alpha_{tCP}) \geq R_m^{OPT}/(c^{max}\alpha_{tCP})$. Second, to show the revenue bound of segments in Ψ_2, let $\hat{R}_{m^*} = \max\{\hat{R}_{m'} | \Phi_{m'} \in S^{OPT} \setminus \Psi_1\}$. For each $\Phi_m \in \Psi_2$, since Athena chooses Φ_m rather than Φ_{m^*} as winning segment, $\hat{R}_m \geq \hat{R}_{m^*}$. So by (10) and [11], we have

$$R_m \geq \hat{R}_m \geq \hat{R}_{m^*} \geq R_{m^*}/(\alpha_{tCP}c_{m^*}) \geq R_{m^*}^{OPT}/(\alpha_{tCP}(c_{m^*})^2).$$

Athena's revenue R can be computed as:

$$\begin{aligned} R &= \sum_{\Phi_m \in \Psi_1} R_m + \sum_{\Phi_m \in \Psi_2} R_m \\ &\geq \sum_{\Phi_m \in \Psi_1} R_m^{OPT}/(c^{max}\alpha_{tCP}) + |\Psi_2| \cdot R_{m^*}^{OPT}/(\alpha_{tCP}(c_{m^*})^2) \\ &\geq \sum_{m \in S^{OPT}} R_m^{OPT}/((c^{max})^2\alpha_{tCP}) = R^{OPT}/((c^{max})^2\alpha_{tCP}), \end{aligned}$$

where $R^{OPT} = \sum_{m \in S^{OPT}} R_m^{OPT}$ is the revenue of OPT that applies single price scheme in each segment in S^{OPT}. □