Tackling Bidder Collusion in Dynamic Spectrum Auctions (Extended)

Xia Zhou, Alessandra Sala, and Haitao Zheng
Department of Computer Science, University of California, Santa Barbara, CA 93106
xiazhou,alessandra,htzheng@cs.ucsb.edu

Abstract—**Dynamic spectrum auction is an effective solution to manage spectrum across many small networks. As the number of participants grows, collusion poses a serious threat to auction performance. Small groups of colluding bidders can make use of the interference constraints to manipulate auction outcomes, leading to unfair spectrum distribution and significant loss in auction revenue. Prior designs, however, are either forced to give up spatial reuse for collusion-resistance, become computationally prohibitive, or can only address very limited types of collusion. In this paper, we present DC^2, a systematic auction design that can effectively discourage collusion and achieve spatial reuse, even when multiple collusion groups are present. DC^2 achieves this using a novel 3-stage "Divide, Conquer, and Combine" procedure that integrates an efficient spectrum allocation algorithm with a powerful collusion-resistant mechanism design. More importantly, DC^2 can configure the level of collusion-resistance and maximize auction revenue for any given level. Auctioneers can now configure auctions based on their own preferences and deployment environments. We analytically prove DC^2's collusion-resistance and its revenue bound, and perform extensive network simulations to verify DC^2's effectiveness. We show that it is particularly effective against small-size collusion, the most commonly observed in practical auctions.**

I. Introduction

In the last few years we have witnessed the flourish of "small" wireless networks. Individual corporate, campuses and communities are deploying a variety of small to medium-size wireless networks in their local neighborhoods. Most of these networks exploit unlicensed bands to achieve market availability and rapid growth. Now they are suffering excessive interference and poor performance due to aggressive deployment and unprotected spectrum usage.

As small networks continue to grow, a pressing problem is how to provide them with proper spectrum usage. Dynamic spectrum auction [1]–[5] offers an attractive solution. It allows many small players to bid for spectrum by the amount they actually need, and exploits time and spatial reuse to improve allocation efficiency. In practice, commercial spectrum trading/auction systems have already started to serve small network providers in the states such as South Carolina [6].

As many small networks seek to obtain spectrum via auctions, collusion becomes a serious threat to auction revenue and efficiency [7]–[9]. Extensive measurements [10], [11] have shown that in many past auctions including the FCC spectrum auctions, a small fraction (<5%) of bidders have strategically formed one or multiple collusion groups and rig their bids to manipulate auction results, causing lower prices and unfair resource distribution. Because collusion is legally banned in

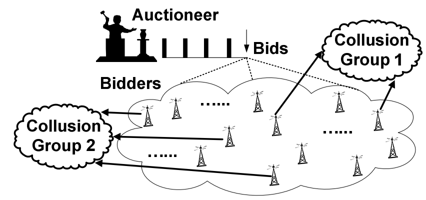

Fig. 1. A large-scale dynamic spectrum auction that contains multiple small-size collusion groups (2–3 players). Because the auction is sealed-bid, bidders have no knowledge of others' bids and behaviors unless they collude. Being rational, one colluding group will only cheat if this improves the total members' utilities.

commercial auctions, existing collusion groups were tacit and small in size (Figure 1), thus easier to form and hard (and expensive) to detect in large-scale auctions. Similar trends were observed in other practical deployments including P2P systems where each collusion group contains 2–4 players [12]. In this paper, we show that in spectrum auctions, small-size collusion is even more effective than that in conventional auctions, because colluding bidders can exploit the bidder interference constraints to rig their bids. These observations pose an important need for spectrum auctions to resist bidder collusion, particularly small-size collusion.

While prior works have proposed solutions to tackle collusion in auctions, we show that, when applied to spectrum auctions, they either cause severe interference or lose collusion-resistance [13]. This is because conventional designs [14], [15] do not consider any reuse and assume bidders have a homogeneous relationship: either all conflict with each other, or do not conflict at all. In spectrum auctions, the relationship becomes heterogeneous due to the bidder interference constraints and the need for spatial reuse [5]. These solutions either lead to heavy interference, or lose collusion-resistance. On the other hand, recent work on spectrum auctions focuses on suppressing some forms of collusion [16], but can be attacked easily by other simple forms. More importantly, this solution requires an exponential-complexity algorithm to ensure its resistance, thus cannot operate in large-scale dynamic spectrum auctions.

In this paper, we propose DC^2, a new collusion-resistant and computationally-efficient spectrum auction. Using a randomization technique, DC^2 resists collusion by diminishing the gain of any colluding group unless it becomes large (and hence hard to form and easy to be detected). Such diminishing returns leave bidders little or no incentive to collude. Meanwhile, DC^2 enables spatial reuse to improve auction revenue and efficiency. DC^2's novel contribution is to let the auctioneer secretely perform a 3-stage "Divide, Conquer, and Combine" procedure after receiving bids. By judiciously designing its procedure, DC^2 successfully integrates an efficient spectrum allocation algorithm (in "Divide") with a novel economic mechanism (in "Conquer"), enabling spatial reuse and effectively controlling collusion.

DC^2 achieves the following key advantages:

- To our best knowledge, DC^2 is the first solution to tackle bidder collusion in large-scale spectrum auctions. It operates in real-time with polynomial complexity, and enables spatial reuse to better distribute the spectrum.
- DC^2 implements a soft and customizable collusion-resistance, referred to as the (t, p)-truthfulness. It ensures that with a probability of p or higher, no colluding group of size t or less can gain any benefit even if multiple colluding groups are present.
- A key component of DC^2 is to configure the auction to maximize its revenue while guaranteeing the required resistance. In particular, we prove that DC^2's revenue is within a factor of $c^{max}\alpha^l_{tCP}$ from the optimum defined by the underlying spectrum allocation. The revenue scales gracefully with the resistance level (t, p), and remains significantly higher than that of the prior solution.
- In designing DC^2, we identify a fundamental tradeoff between collusion-resistance and auction revenue. To resist collusion, an auction must process bids "with caution." In cases where no bidders cheat, this leads to some loss in the revenue compared to the designs without collusion-resistance. By adjusting (t, p), DC^2 can flexibly customize the auction based on the auctioneer's resistance preference and the amount of the revenue she is willing to sacrifice for such robustness.
- DC^2 takes as input any reusability-driven spectrum allocation algorithm, thus applies to both graph and physical interference models.

After analytically proving DC^2's collusion-resistance and revenue bound, we also perform network simulations to verify DC^2's performance and its dependency on the spectrum allocation algorithm. Our results lead to several key findings:

- DC^2 outperforms existing collusion-resistant solution by up to 50% in average revenue when resisting small-size collusion in a large-scale network.
- DC^2 is the most effective against small-size collusion, incurring little cost in revenue. In our experiments, DC2 resist any collusion group no larger than t=2, 4, 8 with $p = 0.8$ by sacrificing only 5%, 14%, 21% revenue compared to the design without collusion-resistance [5]. This applies even when multiple collusion groups exist as long as each group has t bidders or less.
- DC^2 can use any allocation algorithm to form submarkets, but favors the partitions producing large submarkets. Existing spectrum allocation algorithms perform similarly in DC^2 because they generate similar-sized large independent-sets and differ only in small or medium-sized sets. Finding the optimal spectrum allocation algorithm in DC^2 is an interesting open research.

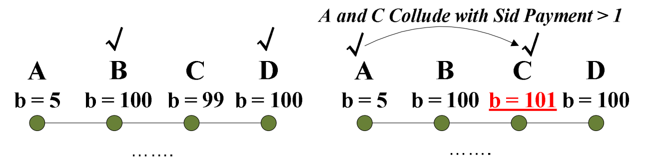

Fig. 2. A collusion example in VERITAS [5] using a section of a *large* network (represented by the conflict graph [21]). (Left) With no collusion, B and D are winners and the auction revenue is $99 \times 2 = 198$. (Right) When A and C collude, A and C are the winners, reducing the revenue to 100.

II. Collusion in Spectrum Auctions

Consider the following dynamic spectrum auction proposed in [5]. An auctioneer runs an auction periodically. Each time it auctions off K channels to some $n(n >> 1)$ bidders who submit bids privately. Each bidder requests one channel and treats the K channels homogeneously. Each bidder holds a valuation of the channels from either economic modeling [17], [18] or network survey. After receiving bids, the auctioneer determines winners and their spectrum usage based on the (complex) interference constraints among the bidders. An important requirement is to enable spatial reuse where non-conflicting bidders can reuse the same channel. Such requirement makes dynamic spectrum auctions fundamentally different from conventional auctions, and imposes significant design challenges. Additional descriptions on dynamic spectrum auctions can be found in [5].

Collusion occurs in an auction when groups of bidders coordinate their bids to game the system, gaining unfair advantages and harming others. Multiple collusion groups might appear in a single auction, but each collusion group is rational [14]. With no knowledge on other bidders' bids and behaviors (because the auctions are sealed-bid), a collusion group will only rig their bids if this can improve the group utility. The group utility refers to the sum of each member's utility, which is defined as its valuation minus its price paid for being a winner, otherwise 0. This notion is stronger than the *group strategyproofness* [19], [20] which assumes the profit cannot be transferred within one collusion group.

Small-Size Collusion is the Bottleneck. There have been several measurement studies and empirical analysis on collusion in past auctions, including spectrum auctions. One key observation is that individual collusion groups are small but cause significant damages. For example, an empirical analysis on four largest FCC spectrum auctions suggests that only a small fraction (1–2%) of bidders have colluded [10]. In the well-known PCS spectrum auction [11], the total number of colluders is no more than 6 out of 153 bidders, yet they won more than 40% of the spectrum auctioned and paid significantly less. Finally, such observation also applies to other non-auction systems. In a commercially deployed P2P system with 160,000+ participants, the dominant collusion groups were very small (2–4 players per group), but highly effective [12].

Another key observation made by our own analysis is that small-size collusion is even more effective in dynamic spectrum auctions. The heterogeneous interference constraints open up new vulnerabilities to collusion, particularly because varying one bidder's bid can create "chain effects" and affect auction results at many other bidders. Consider an example in Figure 2 that uses the truthful auction design from [5], represented by the bidder conflict graph [21]. Assume bidders compete for one channel, $K = 1$. When everyone bids truthfully, B and D are winners, each pays C's bid (99), and the revenue is $99 \times 2 = 198$. If C colludes with A and bids 101, then C and A become winners, paying 100 and 0 respectively. Now A has a payoff of 5 and can compensate C's overbidding via side payment of 2. They gain an unfair advantage and cut the auction revenue to nearly half (from 198 to 100).

This above collusion is effective because of the heterogeneous interference constraints: A can win with low price because its colluder C blocks its competing peer B. In this way, colluding bidders can utilize this spatial reuse property to manipulate auction outcomes, although individually they cannot (since the auction is truthful [5]).

III. Why Existing Solutions Fail?

After identifying the characteristics of collusion in spectrum auctions, we now examine existing solutions that tackle collusion and their effectiveness in spectrum auctions. We categorize them by the types of collusion they address.

A. Solutions That Tackle Some Forms of Collusion

A recent work [16] resists two specific forms of collusion in spectrum auctions. The authors first convert the problem to finding the maximal weighted independent set (MWIS) in the conflict graph weighted by the bids, and then apply the Nash bargaining solution [22] to determine prices. Relying on an exponential complexity algorithm, this solution is only applicable in auctions with a small number of bidders.

In addition, this solution only addresses two specific types of collusion (loser collusion and winner sublease). We show that it is highly vulnerable to a simple winner-loser collusion. Consider an example in Figure 3 where A–C are a segment of a large bidder network, and $K = 1$. When everyone bids truthfully, A and C are winners, paying 100 in total(1.5 and 98.5 respectively). C's payoff is 0.5. Now if C colludes with B and lowers its bid to let B win, B only pays 2.1. So B can split any amount ≥ 0.5 from its surplus to attract C to join. Yet the auction revenue reduces drastically from 100 to 2.1.

The above example demonstrates the complex behavior of collusion. Under the complex interference constraints, it is unrealistic to predict all collusion behaviors. Therefore, it is desirable to design solutions that can address any form of collusion in large-scale auction systems.

B. Solutions That Tackle All Forms of Collusion

It has been proven that completely eliminating the impact of all-size collusion is extremely hard, and the only solution is a trivial "posted price" mechanism [14]. The auctioneer sets a clearing price independent of the bids and those bidding no less than the clearing price are auction winners. This solution, however, leads to unbounded loss in auction revenue.

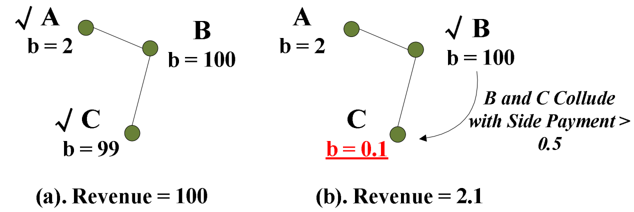

Fig. 3. A collusion example in an existing solution [16]. (Left) Without collusion, A and C are winners and the auction revenue is 100. (Right) B and C collude, making B the winner and reducing the auction revenue to 2.1. B pays C on the side, giving C a higher payoff.

Others [14], [15] consider the notion of *soft collusion-resistance*. They propose to set the price such that when a subset of t bidders collude, the auction outcome is unlikely to change [14], or gives colluders limited gain [15]. Existing solutions in this category, however, were developed for conventional auctions without considering any reusability. They assume that bidders either all conflict with each other [15], or do not conflict at all [14]. When applied to spectrum auctions, they either suffer severe interference or lose collusion-resistance. For example, one simple extension is to first find a feasible price region that leads to interference-free spectrum allocation [3], then apply the algorithms in [14], [15] to set a uniform price. This extension, however, is not collusion-resistant since bidders can rig the bids to control the feasible price region and hence the auction outcome. In addition, with spatial reuse, uniform pricing suffers significant loss in auction revenue as shown in [3].

In summary, existing solutions on collusion-resistant spectrum auctions either cannot address all forms of collusion, require exponential complexity, or must give up spatial reuse completely for collusion-resistance. Thus, there is no practical solution for spectrum auctions that can resist any type of, especially small-size collusion when a large number of bidders are present. This motivates us to identify new auction designs.

IV. DC^2: Divide, Conquer, and Combine

We introduce DC^2, a new auction design that resists *any* form of collusion while enabling spatial reuse. DC^2 applies the concept of soft collusion-resistance [14]– diminishing the gain of any collusion group by making the auction outcomes "insensitive" to bid changes of the group.

Putting it more formally, consider n bidders among which π is one of the colluding groups and has t bidders. In π's view, if $\mathbb{B}$ is the bids of all n bidders when π does not cheat and $\mathbb{B}'$ is the bids when π cheats, then $\mathbb{B}$ and $\mathbb{B}'$ differ by no more than t bids. Being rational, π will have little incentive to cheat if with a probability of p or higher, the auction procedure f_{DC^2} returns the same auction price and result Γ over $\mathbb{B}$ and $\mathbb{B}'$:

$$f_{DC^2}(\mathbb{B}) = f_{DC^2}(\mathbb{B}') = \Gamma. \quad (1)$$

Thus to resist all forms of collusion groups each of size t or less, an auction design needs to ensure that with a high probability, any two sets of bids $\mathbb{B}$ and $\mathbb{B}'$ which differ in no more than t bids will lead to the same auction result [14], [23]. The corresponding soft collusion-resistance is defined as the (t, p)-truthfulness [14], [23]:

***Definition** 1: An auction achieves the (t, p)**-truthfulness** if with a probability of **no less than** p, no collusion group of size t or less can improve its group utility by rigging the bids.*

Note that $(t = 1, p = 1)$ refers to the truthfulness achieved by [5], and $(t \geq 2, p = 1)$ maps to the *hard collusion-resistance* only achievable using the posted price [14]. Thus, DC^2 focuses on $(t \geq 2, p < 1)$.

While prior works have enforced soft collusion-resistance in conventional auctions, we have shown that they fail when applied to spectrum auctions with spatial reuse [13]. DC^2 overcome this challenge using the concept of "Divide and Conquer," integrating a spectrum allocation with a collusion-resistant design. In DC^2 the auctioneer secretely performs the following steps after collecting bids:

- "**Divide**" – The auctioneer applies a spectrum allocation algorithm to divide the bidders independent of their bids into several non-overlapping *sub-markets* such that in each sub-market no bidders interfere with each other.
- "**Conquer**" – In each sub-market, the auctioneer applies a classical collusion-resistant mechanism to identify potential winners, and treats them together as a *super bidder* representing the current sub-market.
- "**Combine**" – Given the available channels, the auctioneer selects winning sub-markets, in which the potential winners become final winners. This stage is essential to enforce spectrum reuse and maintain collusion-resistance.

The key challenge facing DC^2 is to design the integration judiciously to ensure collusion-resistance, because now colluders can manipulate their bids to affect not only the auction result in each sub-market, but also the sub-markets they are assigned to. In the following we discuss the detailed DC^2 design.

A. Divide: Forming Sub-markets

DC^2 first divides bidders *independent of their bids* into several non-overlapping sub-markets based on their interference conditions. Bidders in each sub-market do not conflict with each other and can reuse the same channel. This is done by using any spectrum allocation algorithm to *virtually* assign one channel per bidder without limiting the number of channels used. The channel allocation must be bid-independent, otherwise bidders can rig bids to affect the set of sub-markets formed. Assume V virtual channels are used to satisfy every bidder. Then those allocated with the same channel are grouped into one sub-market, hence forming V sub-markets: $\Phi_1, ..., \Phi_m, ..., \Phi_V$. Table 1 summarizes the procedure.

Table 1. Forming Sub-Markets in DC^2

STEP 1	Apply a spectrum allocation algorithm to assign each bidder with one virtual channel, *independent of the bids*.
STEP 2	Group bidders with the same channel into a sub-market Φ_m.
STEP 3	Return sub-markets: $\Phi_1, ..., \Phi_m, ..., \Phi_V$.

Note that DC^2 can accommodate various interference models since it can use any spectrum allocation algorithms such as [4], [24], [25] to assign one virtual channel per bidder.

B. Conquer: Virtual Clearing

Now each sub-market is free of interference constraints, DC^2 diminishes the impact of a colluding group by applying a *virtual clearing* process in each sub-market. The methodology in this stage is to virtually run an auction in each sub-market to determine its winners and prices, and then represent each sub-market by a *super bidder* who will compete in the "combine" stage. In addition to making each virtual auction collusion-resistant, DC^2 judiciously designs the bid for each super bidder to ensure that it is insensitive to changes of collusive bids. After a proper integration in the "combine" stage, this translates into the collusion-resistance in the entire auction.

With bidders in each sub-market free of interference constraint, enforcing collusion-resistance in each sub-market is much simpler. Well-developed solutions already exist. In this paper, we use the t-Truthful with Probability (tCP) solution in [14]. We will explain tCP in detail in Section V. By applying tCP in Φ_m, DC^2 sets a *virtual price* Γ_m which is insensitive to collusive bids, and marks $N(\Gamma_m)$ potential winners as those bidding no less than Γ_m. Treating Φ_m as a super-bidder, DC^2 computes an estimated revenue $\hat{R}(\Gamma_m)$ as its bid:

$$\hat{R}(\Gamma_m) = \Gamma_m \times g_c(N(\Gamma_m)), \tag{2}$$

where $g_c(\cdot)$ is a random rounding function that makes $\hat{R}(\Gamma_m)$ insensitive to $N(\Gamma_m)$ which could be affected by collusive bids. This procedure is the key to ensure collusion-resistance of the integration. Table 2 summarizes the actions in this stage.

Table 2. Virtual Clearing in Sub-Market Φ_m

STEP 4	1) Apply tCP in Φ_m to set the virtual price Γ_m; 2) Compute an estimated revenue $\hat{R}(\Gamma_m)$.
STEP 5	1) Identify the potential winners W^{Φ_m} as those with bids no less than Γ_m; 2) Represent W^{Φ_m} as a super bidder with bid $\hat{R}(\Gamma_m)$.
STEP 6	Return $\hat{R}(\Gamma_m)$, W^{Φ_m}, and Γ_m.

C. Combine: Final Clearing

In this stage, the auctioneer selects the final winners by examining the V super bidders. Given K available channels for auction, the auctioneer will choose $\min(K, V)$ super bidders (sub-markets) with the highest bids as winners and assign one channel per winning sub-market. Each potential winner in each winning sub-market becomes the final auction winner, and is charged by the virtual price in its sub-market determined in the "conquer" stage. In this way, winners in the same sub-market are charged equally, but those from different winning sub-markets could be charged differently. Note that we cannot recycle the losers of one winning sub-market to another, because it breaks the requirement of bid-independent sub-market formation that is essential to ensure collusion-resistance. We summarize the actions in Table 3.

DC^2 Optimization. A key component of DC^2 is to configure its auction procedure to maximize the auction revenue while guaranteeing the required level of collusion-resistance defined

Table 3. Final Clearing in DC^2

STEP 7	Choose $\min(K, V)$ highest super bidders (sub-markets) as winners; assign one channel to each.
STEP 8	In each winning sub-market Φ_m, its virtual winners W^{Φ_m} are real winners; each gets a channel and is charged by Γ_m.

by the (t, p)-truthfulness. In the following two sections, we first analytically prove DC^2's collusion-resistance, and then describe DC^2's detailed configuration to maximize auction revenue for a given (t, p). We also explore their dependency on the configuration of sub-markets.

V. DC^2's Collusion Resistance Analysis

This section analyzes DC^2's collusion-resistance, and its dependency on sub-market sizes. As discussed in Section IV, we measure the level of collusion-resistance by the (t, p)-*truthfulness*. Our main result is:

***Theorem** 1: DC^2 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 sub-market that uses tCP for virtual clearing, c_{min} and λ are auction parameters (defined in Section V-B). When $t/l_{min} << 1$, we have $p = 1 - O(t/l_{min})$.*

DC^2 achieves this resistance by integrating a reusability-driven spectrum allocation algorithm with a collusion-resistant tCP mechanism. While tCP only applies to auctions with a single sub-market, DC^2 judiciously configures the auction to achieve collusion-resistance in the presence of multiple sub-markets. Next we first provide some background on tCP, and then present the proof of Theorem 1.

A. Preliminary

As background, we briefly describe tCP and show that for each sub-market it achieves the (t, p)-truthfulness. Consider an auction that contains just one sub-market Φ_m with a bid set $\mathbb{B}_m$ where bidders do not conflict with each other. 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:

- Choose a parameter $\alpha > 1$;
- Define $\mathbb{G} = \{\alpha^i | i \in \mathbb{Z}\}$ 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 $\mathbb{B}_m$;
- Use a *consensus estimation* function $g_c(\cdot)$ parameterized by c to randomly round $N(\alpha^i)$; $g_c(\cdot)$ 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 [26];
- Compute $\Gamma_m = \arg\max_{\alpha^i \in \mathbb{G}} \alpha^i \cdot g_c(N(\alpha^i))$.

The random rounding $g_c(\cdot)$ makes price selection insensitive to bid changes, ensuring that with high probability, Γ_m will not be affected by no more than t bids. As proved in [14], tCP achieves the collusion-resistance of the (t, p)-truthfulness, where p is lower bounded as 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.*

This bound indicates that tCP favors cases with large number of winners $l >> t$. In these cases, we have $p = 1 - O(t/l)$.

B. Proof of Theorem 1

We consider the scenario where the auction contains multiple sub-markets, since Lemma 1 directly applies when there is a single sub-market.

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

Next in "Conquer", assume there are V sub-markets where each Φ_m uses tCP with (c_m, α_m) for virtual clearing and has l_m winners ($m = 1, ..., V$). Let $c^{min} = \min\{c_1, ..., c_V\}$, $c^{max} = \max\{c_1, ..., c_V\}$, $\alpha^{min} = \min\{\alpha_1, ..., \alpha_V\}$, and $l^{min} = \min\{l_1, ..., l_V\}$. Consider any collusion group of size $\leq t$. Let t_m be the number of its members assigned to sub-market 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. From Lemma 1 and let $\lambda_m = \frac{2c_m\alpha_m}{\alpha_m-1}$, we have:

$$Pr_m^{tT} \geq 1 + \log_{c_m}(1 - \frac{\lambda_m t_m}{l_m - t_m}), \tag{3}$$

$$\widetilde{Pr}_m^{tT} = 1 - Pr_m^{tT} \leq -\log_{c_m}(1 - \lambda_m \frac{t_m}{l_m - t_m}). \tag{4}$$

By DC^2's "Combine", if each sub-market 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 DC^2's auction result. Let Pr^{tT} be the probability that DC^2 is $t-$truthful, we have

$$\begin{aligned} Pr^{tT} &\geq \prod_{m=1}^{V} Pr_m^{tT} = \prod_{m=1}^{V} (1 - \widetilde{Pr}_m^{tT}) \geq 1 - \sum_{m=1}^{V} \widetilde{Pr}_m^{tT} \\ \text{(from (4))} \quad &\geq 1 + \sum_{m=1}^{V} \log_{c_m}(1 - \lambda_m \frac{t_m}{l_m - t_m}). \end{aligned} \tag{5}$$

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

$$\begin{aligned} Pr^{tT} &\geq 1 + \sum_{m=1}^{V} \log_{c^{min}}(1 - \lambda_m \frac{t_m}{l_m - t_m}) \\ &= 1 + \log_{c^{min}}(\prod_{m=1}^{V}(1 - \lambda_m \frac{t_m}{l_m - t_m})) \\ &\geq 1 + \log_{c^{min}}(1 - \sum_{m=1}^{V} \lambda_m \frac{t_m}{l_m - t_m}). \end{aligned} \tag{6}$$

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 (6) into

$$\begin{aligned} Pr^{tT} &\geq 1 + \log_{c^{min}}(1 - \lambda \frac{\sum_{m=1}^{V} t_m}{l^{min} - t}) \\ &\geq 1 + \log_{c^{min}}(1 - \lambda \frac{t}{l^{min} - t}). \end{aligned} \tag{7}$$

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 $p = 1 - \mathrm{O}(t/l^{min})$. ■

C. Impact of Sub-Market Configuration

From Theorem 1 we see that l^{min} is critical to DC^2's collusion-resistance. Given (t,p), we derive the minimum l^{min} required to perform tCP on all the sub-markets. Since $\lambda = \frac{2c^{max}\alpha^{min}}{\alpha^{min}-1}$, by Theorem 1 we can derive

$$\alpha^{min} = \frac{(1-(c^{min})^{p-1})(l^{min}-t)}{(1-(c^{min})^{p-1})(l^{min}-t) - 2c^{max}t}. \tag{8}$$

Running tCP requires $\alpha^{min} > 1$ (required by the consensus estimation function [26]), which translates into

$$l^{min} > \frac{2c^{max}t}{1-(c^{min})^{p-1}} + t \triangleq l_{tCP}. \tag{9}$$

In other words, tCP is applicable to all the sub-markets only when there are enough winners in each, which also means that all the sub-markets need to be large enough. This condition l_{tCP} elevates as t, p increase.

The sub-market configuration, on the other hand, depends on the bidder interference constraints and the spectrum allocation algorithm. In many cases, there will be some small sub-markets that do not satisfy the l_{tcp} condition. In this case, DC^2 applies tCP to only large enough sub-markets, and the posted price to others. Because the posted price mechanism achieves hard collusion-resistance, DC^2's collusion-resistance (t,p) only depends on the sub-markets that run tCP.

From Theorem 1 we also see that the choice of l_{min} and the set of sub-markets performing tCP will affect the configuration of λ, c^{min} and thus p. As we will show next, the auction revenue also depends on these parameters. Therefore, an important challenge in DC^2 is to configure the auction carefully to achieve a given (t,p)-truthfulness and at the same time maximize the auction revenue.

VI. (t,p)-BASED REVENUE MAXIMIZATION

In this section, we discuss how to configure DC^2 to maximize the auction revenue while achieving a given (t,p) requirement. After forming a set of sub-markets via a spectrum allocation algorithm, DC^2 needs to determine the set of sub-markets that will run tCP as virtual clearing, and their configurations c_m, α_m. Note that to ensure collusion-resistance, the sub-market formation and configuration must be independent of the bids. Overall, our proposed configuration has the following analytical guarantee on revenue:

***Theorem** 2: While satisfying the (t,p)-truthfulness, DC^2 with V sub-markets running tCP achieves an auction revenue no less than $R^{OPT}/(c^{max}\alpha^l_{tCP})$, where R^{OPT} is the sum of the optimal revenue obtained by treating each of the V sub-markets separately, and c^{max}, α^l_{tCP} (defined by (12), (13)) are auction parameters required to achieve the (t,p)-truthfulness.*

In this following, we first explain the factors that affect DC^2's revenue, the challenges and our solutions. We then prove Theorem 2 and examine DC^2's overall complexity.

A. Factors that Affect the Revenue

We show that the auction revenue depends on two factors, from which we can derive the relationship among l^{min}, c^{max}, α^l_{tCP} for a given (t,p).

(1) Each sub-market's virtual clearing mechanism. Each sub-market would prefer tCP over posted-price as its virtual clearing mechanism. This is because in the context of a single sub-market, tCP with (c_m, α_m) can guarantee that the revenue is within a distance of $c_m\alpha_m$ from the optimal [26], while posted-price provides no guarantee. However, as shown in Section V-C, only large sub-markets can use tCP in order to satisfy the (t,p) requirement. The larger the t and p, the fewer number of sub-markets can run tCP, and the less the revenue.

(2) The choice of (c_m, α_m) in a sub-market that runs tCP. Ideally, c_m and α_m should be set as small as possible to maximize the revenue produced from tCP. If they are set improperly, the expected revenue could be lower than that of running the posted price. In particular, as we will show in Section VI-C, α_m must be small enough so that the average revenue of tCP is higher than that of posted-price. So when considering revenue, a sub-market Φ_m runs tCP if,

$$\alpha_m \le \alpha^u_{tCP}. \tag{10}$$

On the other hand, c_m and α_m must be large enough to achieve the required collusion-resistance. From [26], we have

$$c_m = \mathrm{argmax}_x[(\frac{l_m - t}{l_m + t} - \frac{1}{x})/\ln(x)], \tag{11}$$

$$c^{max} = \max_{m, l_m \ge l^{min}} c_m, \quad c^{min} = \min_{m, l_m \ge l^{min}} c_m. \tag{12}$$

From (8) in Section V-C, we have

$$\begin{aligned} \alpha_m \quad &\ge \quad \alpha^{min} = \frac{(1-(c^{min})^{p-1})(l^{min}-t)}{(1-(c^{min})^{p-1})(l^{min}-t) - 2c^{max}t} \\ &\triangleq \quad \alpha^l_{tCP}. \end{aligned} \tag{13}$$

B. Challenges in Configuring Revenue-Maximizing DC^2

The above analysis shows that having multiple sub-markets brings significant challenges in DC^2. First, the configuration parameters are inter-dependent. From the above, a sub-market Φ_m should run tCP if

$$\begin{aligned} l_m \ge \quad & l^{min} \quad \ge l_{tCP}, \\ \alpha^l_{tCP} \le \quad & \alpha_m \quad \le \alpha^u_{tCP}. \end{aligned}$$

We can set $\alpha_m = \alpha^l_{tCP}$ to maximize auction revenue. The choice of l_{tCP} and α^l_{tCP}, however, depends on which and how many sub-markets choose to run tCP, *i.e.* l^{min}. Such inter-dependency creates a dilemma in choosing each sub-market's virtual clearing mechanism. Intuitively, the system should let as many sub-markets run tCP as possible. But as more smaller sub-markets start to use tCP, α^l_{tCP} increases and so does each sub-market's $\alpha_m = \alpha^l_{tCP}$. Higher α_m maps to more degradation in each sub-market's revenue and thus the overall revenue. Therefore, we need a mechanism to judiciously choose the virtual clearing mechanism for each sub-market, which we will describe in SectionVI-C.

Second, such inter-dependency also creates a dilemma in forming sub-markets. Creating balanced sub-market partition will allow more sub-markets to apply tCP but each sub-market has a larger α_m and hence less revenue. On the other hand, imbalanced partition could prevent some small sub-markets from applying tCP and also lead to revenue loss. By examining several allocation algorithms, we will study the impact of sub-market formation in Section VII-C.

C. Revenue-Maximizing DC^2

We now describe the revenue-maximizing DC^{2}design, for a given (t, p) and a set of sub-markets formed in the "Divide" stage. From the above descriptions, we see that the configuration reduces to judiciously setting parameters $(l_{tCP}, \alpha^l_{tCP})$, such that only sub-markets with $\geq l_{tCP}$ winners will run tCP with operating parameter α^l_{tCP} and others run posted price. DC^2 chooses the best configuration that achieves (t, p)-truthfulness and maximizes the revenue.

To maintain collusion-resistance, sub-market configuration must be independent of the bids. Thus when evaluating each candidate configuration, the system cannot use the actual bids to compute the auction revenue. Instead, DC^2 uses a statistical method to estimate the revenue assuming the bids follow a predefined random distribution. For example, with no knowledge on the actual bid distribution, DC^2 can use the uniform distribution. In this case, we can derive the expected revenue of a sub-market with x bidders when it runs tCP ($E^{tCP}_{\alpha}(x) = \alpha^{-1}(1 - \alpha^{-1})x$) and when it runs posted price ($E^P(x) = x/6$). We also show that $\alpha^u_{tCP} = (3 + \sqrt{3})$. We list the detailed derivations in Appendix and the overall procedure for setting $(l_{tCP}, \alpha^l_{tCP})$ in Algorithm 1.

It should be noted that in its basic form, DC^2 does not require any knowledge on the bid distribution. When such information is available, DC^2 can use this knowledge to improve the revenue estimation. In Section VII, we will evaluate the performance of DC^2 under different bid distributions.

Complexity Analysis. DC^2's complexity comes from the 3-stage procedure described in Section IV, and the above configuration Algorithm 1. Consider n bidders that form V sub-markets and bid for K channels, and assume $V \geq K$. In the main procedure, 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 bids $O(n)$; and the Combine stage takes $O(V \log(V))$ time to sort super bidders. In the configuration algorithm, it takes $O(V \log(V))$ to sort sub-markets' sizes, and $O(V)$ to find the statistically optimal parameter configuration. Clearly $V \leq n$, hence the overall complexity of DC^2 is $O(n \log(n))$, plus the complexity of the spectrum allocation algorithm.

D. Analysis of Revenue Bound

We now present the proof of Theorem 2.

Proof: Let $c_1, ..., c_V$ be the auction parameter for each sub-market defined in (11). We can then compute c^{min}, c^{max}

Algorithm 1 DC^2-Configuration(t, p, $\mathbb{N}$, Y)

$Y = \min(V, K)$ given V sub-markets and K channels;
$\mathbb{N} : N_1 \geq N_2 ... \geq N_Y$, Y largest sub-markets each with $N_i (i \leq Y)$ bidders.

for $m = 1$ to Y **do**
  $l^{min} \leftarrow l_m = \lceil \frac{N_m}{2} \rceil$
  $c^{max} \leftarrow c_m$ by (11)
  $c^{min} \leftarrow c_1$
  if no enough estimated winners by (9) **then**
    $E(m) = \sum_{i=1}^{Y} E^P(N_i)$
    $c_m = 0$
  else
    $a(m) \leftarrow \alpha^{min}$ by (13)
    if $\alpha^{min} \geq \alpha^u_{tCP}$ **then**
      $E(m) = \sum_{i=1}^{Y} E^P(N_i)$
      $c_m = 0$
    else
      $E(m) = \sum_{i=1}^{m} E^{tCP}_{\alpha^{min}}(N_i) + \sum_{i=m+1}^{Y} E^P(N_i)$
    end if
  end if
end for
$m^* = \text{argmax}_m E(m)$
if $c_{m^*} > 0$ **then**
  Return $(l_{tCP}, \alpha^l_{tCP}) = (l_{m^*}, a(m^*))$
else
  Return $(0, 0)$
end if

and α^l_{tCP} based on (t, p). In each sub-market Φ_m, let R_m denote the revenue achieved by DC^2, and R^{OPT}_m be the optimal revenue that can be achieved by a single-price auction. Since $\alpha_m = \alpha^l_{tCP}$ (Section VI-B), by tCP's property [14], we have for each $m \leq V$,

$$R_m \geq R^{OPT}_m / (c_m \alpha^l_{tCP}) \geq R^{OPT}_m / (c^{max} \alpha^l_{tCP}).$$

The overall DC^2's revenue R is

$$R = \sum_{m=1}^{V} R_m \geq \sum_{m=1}^{V} R^{OPT}_m / (c^{max} \alpha^l_{tCP}) = R^{OPT} / c^{max} \alpha^l_{tCP},$$

where $R^{OPT} = \sum_{m=1}^{V} R^{OPT}_m$ is the sum of the optimal revenue that can be achieved by applying a single price mechanism in each sub-market. ■

VII. Evaluation

In this section, we perform network simulations to evaluate DC^2. We identify the tradeoff between collusion-resistance and auction revenue by comparing DC^2 to existing solutions. We then evaluate DC^2 under inaccurate estimation of bid distribution. Finally, we examine the impact of sub-market partition using various spectrum allocation algorithms.

Using large-scale auction systems, we compare the following three solutions with different levels of collusion-resistance:

- **Posted Price [14]**: the only solution achieving hard resilience to collusion ($t = n, p = 1$). It picks a price randomly independent of bids;
- **DC^2**: our proposed solution, providing a soft (t, p) collusion-resistance, where t is the maximal per-group size of all collusion groups;
- **VERITAS [5]**: a truthful spectrum auction design that cannot address collusion.

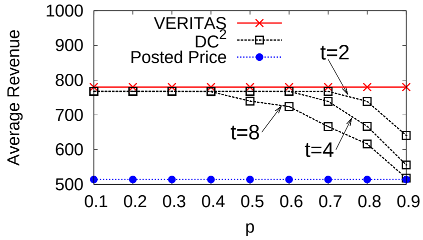

(a) Average revenue vs. p when the maximal collusion group size $\leq 2, 4, 8$.

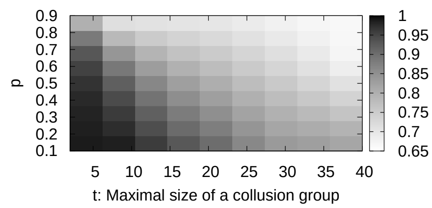

(b) DC^2's normalized revenue indicated by the darkness as a function of (t, p).

Fig. 4. Tradeoff between resilience and revenue. (a) compares various designs; (b) shows the ratio of DC^2's revenue over VERITAS at various (t, p).

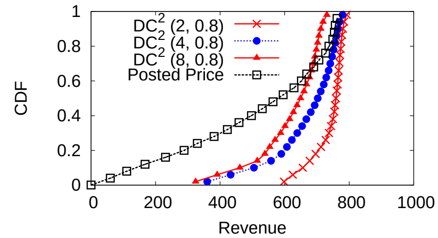

(a) CDF of revenue under uniform bid distribution.

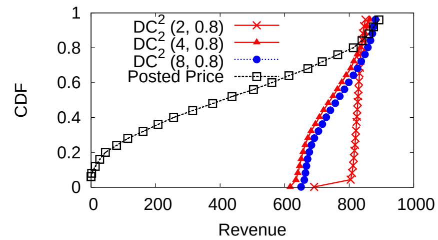

(b) CDF of revenue under non-uniform bid distribution.

Fig. 5. Evaluating DC^2 under various bid distributions by comparing to posted price when each collusion group $\leq 2, 4, 8$, $p = 0.8$.

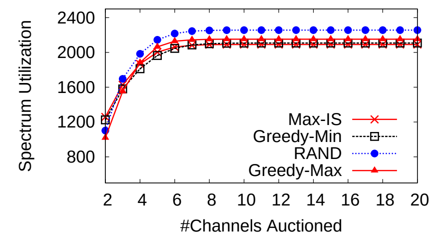

(a) Spectrum utilization

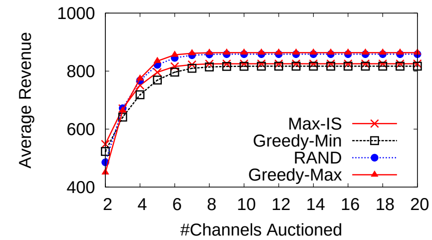

(b) Average revenue

Fig. 6. The performance of DC^2 when applying four different spectrum allocation algorithms. $(t, p) = (4, 0.8)$, where t is the maximal size of one collusion group.

Focusing on large-scale auctions, we did not examine [16] because it only applies to small networks and resists some special forms of collusion. We did not compare soft collusion-resistant solutions in [14], [15] since their allocations are not conflict-free and simple extensions lose collusion-resistance [13].

We simulate a large auction system and deploy 4000 bidders randomly in a 1×1 area. We model the bidder conflict constraints using a distance-based criterion (two bidders within a distance of 0.02 conflict with each other). By default, the bids are randomly distributed in $(0, 1]$, the average conflict degree is roughly 5. We apply a well-known greedy allocation algorithm [24] to form sub-markets. For example, in one topology we obtain 9 sub-markets of sizes 1308, 1025, 737, 479, 277, 114, 40, 15, and 5.

Performance metrics. Both posted price and DC^2 are randomized solutions while VERITAS is a deterministic solution. We compare them in terms of the average revenue (over 5000 bid generations) and the revenue's cumulative distribution function. Moreover, we compare these solutions assuming that everyone bids truthfully and thus the bids are the same for these solutions. In terms of revenue, this is the worst case for DC^2. The difference between DC^2 and VERITAS represents the maximum cost in revenue required to achieve collusion-resistance, while the difference between DC^2 and posted price represents the cost saving by using soft collusion-resistance and focusing on small-size colluding groups.

A. Tradeoff between Robustness and Revenue

Figure 4(a) compares the three solutions in terms of the average revenue, where the revenue of DC^2 depends on the required (t, p). Compared to VERITAS, DC^2 sacrifices 5%, 14% and 21% revenue to achieve collusion-resistance with $t \leq 2, 4, 8$, and $p = 0.8$. This demonstrates DC^2's effectiveness in resisting small-size collusion groups. Compared to posted price, DC^2 improves the revenue by as much as 50% even with $p = 0.8$. As the required collusion-resistance (t, p) gets stronger, the number of sub-markets running tCP reduces and the revenue decreases. Eventually, DC^2 falls back to posted price. Overall, DC^2 introduces collusion-resistance to spectrum auctions at very little overhead, and offers an important flexibility of configuring the auction to exploit the tradeoff between revenue and collusion-resistance.

To further examine the tradeoff, we show in Figure 4(b) the normalized revenue of DC^2 over VERITAS for various (t, p). As expected, DC^2's revenue decreases as (t, p) increases. We see that for $t \leq 5$, the revenue degradation is significantly lower. This again verifies that DC^2 is effective over small-size collusion group, the dominant type in practical auctions.

B. Robustness to Inaccurate Estimation of Bid Distribution

DC^2 runs without requiring accurate information on bid distribution. To examine this robustness, we use a uniform bid distribution function in DC^2's configuration, and generate actual bids by uniform and other non-uniform distributions. Figure 5(a)-(b) compare DC^2(at different(t, p)) to posted price under the same set of bids following the same uniform bid distribution and a Beta distribution $(\alpha = 5, \beta = 5)$ as an example. In both cases, we see that DC^2 maintains a much lower variance in its auction revenue. This is because without considering bids, posted price is more likely to use a price either too high or too low, making the revenue oscillate. In contrast, DC^2 assigns a higher probability to the prices close to the optimal, achieving more stable and higher revenue. This holds even for non-uniform bid distribution. Because as more bids become similar, the estimated revenue of each candidate

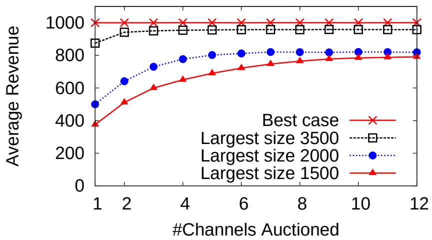

Fig. 7. Comparing various sub-market partitions in DC^2, $(t,p)=(4,0.8)$, where t is the maximal size of one collusion group.

price in tCP differs more. Thus the optimal price is chosen with a higher probability. This demonstrates DC^2's robustness and applicability in practice.

C. Impact of Sub-Market Formation

To understand how the sizes of sub-markets impact the auction revenue and collusion-resistance, we examine four representative spectrum allocation algorithms in DC^2:

- **Max-IS** [25] which assigns channels by finding the maximal independent set of the conflict graph;
- **Greedy-Min** [24] which sequentially assigns channel to the user with the minimal degree in the remaining updated conflict graph. [5] applies it for spectrum allocation;
- **Greedy-Max** which is similar to Greedy-Min except each time choosing the user with the maximal degree in the remaining updated conflict graph;
- **RAND** which picks a user randomly to allocate a channel.

Note that these spectrum allocation algorithms are not designed to optimize DC^2. In principle we have Greedy-MAX $<$ RAND $<$ Greedy-Min $\leq$ Max-IS in terms of spectrum efficiency. Figure 6(a)-(b) compare these solutions in DC^2 by their spectrum utilizations and revenues as the functions of the number of channels. We see that they achieve similar performance. This is because they produce similar-sized large sub-markets, and differ mostly in the forms of small, medium-sized sub-markets. Because DC^2's revenue depends heavily on the sizes of large sub-markets that run tCP, the difference among these algorithms is somewhat diminished.

To further understand the impact of sub-market formation, we now assume that bidders can be partitioned arbitrarily. Figure 7 compares three randomly generated partitions that limit the largest sub-market size to 3500, 2000, 1500 respectively, to the ideal case where all 4000 bidders are in one sub-market. Aside from the ideal case, we see that DC^2 favors the most imbalanced partition (3500 bidders in the largest sub-market), particularly when the number of channels is small. This is because the imbalanced partition leads to larger sub-markets with revenue closer to the optimal one, which compensates having more small sub-markets that run posted price and have no guaranteed revenue. Note that the conventional spectrum allocation algorithms are designed in this manner.

VIII. Related Work

Collusion-resistant design has been widely studied in conventional auctions as well as spectrum auctions. As we have discussed in Section II, [14], [15] have proposed solutions for conventional auctions where bidders conflict with everyone else or do not conflict at all. DC^2 applies the tCP solution [14], but focuses on designing a collusion-resistant solution for dynamic spectrum auctions with general interference constraints. DC^2 also extends tCP to consider combining posted-price and further improve auction revenue. The work in [27] also proposes a pricing game assuming everyone conflicts with each other and hence no spectrum reuse.

Collusion-resistance has been examined in spectrum auctions with spectrum reuse by assuming the specific collusion behaviors [16]. However, in practice the colluding behaviors are highly complex and hard to predict, thus our goal is to provide a general solution that can address any type of collusion. Also, while [16] requires solving NP-hard optimization problems, DC^2 can work with any spectrum allocation solution including polynomial-time ones [4], [24], [25].

IX. Conclusion and Future Work

We propose DC^2, a new spectrum auction design to combat collusion. Using the concept of "Divide and Conquer," DC^2 decouples the problem of spectrum allocation from that of economic mechanism design, achieving spectrum reuse while maintaining collusion-resistance. DC^2 implements a *soft* form of collusion-resistance, enabling the auctioneer to flexibly configure its resilience level and achieve a much better tradeoff between the resilience and the cost in auction revenue.

To our best knowledge, DC^2 is the first to address any form of collusion in spectrum auctions and enable spatial reuse. DC^2 can be extended in several directions. (1) We can consider a different form of soft collusion-resistance by bounding the average collusion gain, which requires a mechanism design different from tCP. (2) Collusion will have more profound impact if bidders can request multiple channels. Addressing collusion in this context requires a stronger rule. (3) DC^2 can use any spectrum allocation algorithm. While the most well-known algorithms perform similarly, it is desirable to find the best allocation algorithm in DC^2 that maximizes the revenue.

Appendix: Deriving $E^P(\cdot)$, $E_\alpha^{tCP}(\cdot)$, and α_{tCP}^u

Assume there are n bidders in the sub-market. Let $F(x)$ be the cumulative distribution of bids in $(0,1]$. In tCP, because the expectation of $g_c(n)$ is $(n/\sqrt{c})$ [26], the expectation of the estimated revenue $\hat{R}(\alpha^i)$ at price α^i is:

$$E(\hat{R}(\alpha^i)) = \alpha^i(1-F(\alpha^i))\frac{n}{\sqrt{c}} \tag{14}$$

We estimate the final price $\Gamma^*(\alpha)$ as the one maximizing $E(\hat{R}(\alpha^i))$, so the expectation of the actual revenue of tCP is

$$E_\alpha^{tCP}(n) = \Gamma^*(\alpha)(1-F(\Gamma^*(\alpha)))n \tag{15}$$

The expected revenue of posted-price E^P is:

$$E^P(n) = n \cdot \int_0^1 \gamma(1-F(\gamma))d\gamma. \tag{16}$$

Thus we can estimate α_{tCP}^u by setting $E^P(n) \geq E_\alpha^{tCP}(n)$.

In the case of uniform bid distribution, $F(x) = x$. So for $\alpha > 2$, $\Gamma^*(\alpha) = \alpha^{-1}$. Then $E_\alpha^{tCP}(n) = \alpha^{-1}(1-\alpha^{-1})n$, $E^P(n) = n \cdot \int_0^1 \gamma(1-\gamma)d\gamma = n/6$. By solving $E^P(n) \geq E_\alpha^{tCP}(n)$:

$$\frac{n}{6} \geq \alpha^{-1}(1-\alpha^{-1})n \tag{17}$$

we have $\alpha \geq (3+\sqrt{3})$, $\alpha_{tCP}^u = (3+\sqrt{3}) \approx 5$.

References

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

[2] O. Ileri, D. Samardzija, and N. B. Mandayam, "Demand responsive pricing and competitive spectrum allocation via a spectrum server," in *Proc. of DySPAN*, 2005.

[3] S. Gandhi, C. Buragohain, L. Cao, H. Zheng, and S. Suri, "A general framework for wireless spectrum auctions," in *Proc. of DySPAN*, 2007.

[4] A. P. Subramanian, M. Al-Ayyoub, H. Gupta, S. R. Das, and M. M. Buddhikot, "Near optimal dynamic spectrum allocation in cellular networks," in *Proc. of DySPAN*, October 2008.

[5] X. Zhou, S. Gandhi, S. Suri, and H. Zheng, "eBay in the sky: Strategy-proof wireless spectrum auctions," in *Proc. of MobiCom*, Sept. 2008.

[6] "Spectrum bridge. http://www.spectrumbridge.com/."

[7] M. Friedman, "Comment on 'collusion in the auction market for treasury bills'," *J. of Political Economy*, vol. 9, pp. 757–785, 1996.

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

[9] P. Cramton and J. Schwartz, "Collusive bidding: Lessons from the fcc spectrum auctions," *Journal of REgulatory Economics*, vol. 17, pp. 229–252, May 2000.

[10] P. Bajari and J. Yeo, "Auction design and tacit collusion in FCC spectrum auctions," National Bureau of Economic Research, Working Paper 14441, October 2008.

[11] P. Cramton and J. Schwartz, "Collusive bidding in the FCC spectrum auctions," University of Maryland, Department of Economics - Peter Cramton, Papers of Peter Cramton, Dec. 2002.

[12] Q. Lian, Z. Zhang, M. Yang, B. Y. Zhao, Y. Dai, and X. Li, "An empirical study of collusion behavior in the maze P2P file-sharing system," in *Proc. of ICDCS*, 2007.

[13] X. Zhou, A. Sala, and H. Zheng, "Tackling bidder collusion in dynamic spectrum auctions (extended)," *UCSB Technical Report*, 2009.

[14] A. V. Goldberg and J. D. Hartline, "Collusion-resistant mechanisms for single-parameter agents," in *Proc. of SODA*, 2005.

[15] F. McSherry and K. Talwar, "Mechanism design via differential privacy," in *Proc. of FOCS*, 2007.

[16] Y. Wu, B. Wang, K. Liu, and T. Clancy, "Collusion-resistant multi-winner spectrum auction for cognitive radio networks," in *Proc. of IEEE Globecom*, Dec. 2008.

[17] P. Bajari and J. T. Fox, "Complementarities and collusion in an FCC spectrum auction," National Bureau of Economic Research, Working Paper 11671, October 2005.

[18] O. Ileri and N. Mandayam, "Dynamic spectrum access models: toward an engineering perspective in the spectrum debate," *Communications Magazine, IEEE*, vol. 46, no. 1, pp. 153–160, January 2008.

[19] K. Jain and V. Vazirani, "Group strategyproofness and no subsidy via lp-duality," in *Proc. of ASPLOS*, 1999.

[20] S. Zhong and F. Wu, "On designing collusion-resistant routing schemes for non-cooperative wireless ad hoc networks," in *Proc. of MobiCom*, 2007.

[21] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu, "Impact of interference on multi-hop wireless network performance," in *Proc. of MobiCom*, 2003.

[22] J. Nash, "The bargaining problem," *Econometrica*, vol. 18, no. 2, pp. 155–162, April 1950.

[23] A. Archer, C. Papadimitriou, K. Talwar, and Éva Tardos, "An approximate truthful mechanism for combinatorial auctions with single parameter agents," in *Proc. of SODA*, 2003.

[24] S. Ramanathan, "A unified framework and algorithm for channel assignment in wireless networks," *Wirel. Netw.*, vol. 5, no. 2, pp. 81–94, 1999.

[25] A. P. Subramanian, H. Gupta, S. R. Das, and M. M. Buddhikot, "Fast spectrum allocation in coordinated dynamic spectrum access based cellular networks," in *Proc. of DySPAN*, November 2007.

[26] A. V. Goldberg and J. D. Hartline, "Competitiveness via consensus," in *Proc. of SODA*, 2003.

[27] Z. Ji and K. Liu, "Multi-stage pricing game for collusion-resistant dynamic spectrum allocation," *Selected Areas in Communications, IEEE Journal on*, vol. 26, no. 1, pp. 182–191, Jan. 2008.