Designing Truthful Spectrum Double Auctions with Local Markets

Wei Wang, *Student Member, IEEE,* Ben Liang, *Senior Member, IEEE,*
and Baochun Li, *Senior Member, IEEE*

Abstract—Market-driven spectrum auctions offer an efficient way to improve spectrum utilization by transferring unused or under-used spectrum from its primary license holder to spectrum-deficient secondary users. Such a spectrum market exhibits strong *locality* in two aspects: 1) that spectrum is a local resource and can only be traded to users within the license area, and 2) that holders can partition the entire license areas and sell any pieces in the market. We design a spectrum double auction that incorporates such locality in spectrum markets, while keeping the auction *economically robust* and computationally efficient. Our designs are tailored to cases with and without the knowledge of bid distributions. Complementary simulation studies show that spectrum utilization can be significantly improved when distribution information is available. Therefore, an auctioneer can start from one design without any *a priori* information, and then switch to the other alternative after accumulating sufficient distribution knowledge. With minor modifications, our designs are also effective for a profit-driven auctioneer aiming to maximize the auction revenue.

Index Terms—Dynamic spectrum access, spectrum double auction, local markets, truthfulness, uniform pricing, discriminatory pricing.

1 Introduction

The recent explosive growth of wireless networks, with their ever-growing demand for radio spectrum, has exacerbated the problem of spectrum scarcity. Such scarcity, however, is not an outcome of exhausted physical spectrum, but a result of inefficient channel utilization due to existing policies that channels are licensed to their authorized holders (typically those who win government auctions of spectrum), and unlicensed access is not allowed even if the channel is not used.

In order to utilize such idle channels and to improve their utilization, it is critical to design sufficient incentives that encourage primary license holders to allow other spectrum-deficient users to access these channels. It is intuitive to observe that under-used channels have values that can be efficiently determined by a *market*, governed by spectrum auctions. If designed well, a spectrum auction offers an efficient way to create a market: it attracts both license holders and wireless users to join, and to either buy or sell idle channels in the market. Once a transaction is conducted, the seller (license holder) earns extra income by leasing unused channels to the buyer (wireless user), who pays to obtain the channel access.

Yet, it is important to point out that transactions take place in secondary markets where spectrum is leased in a seller-defined geographic region. Unlike physical commodities that can be traded anywhere in the world, spectrum is

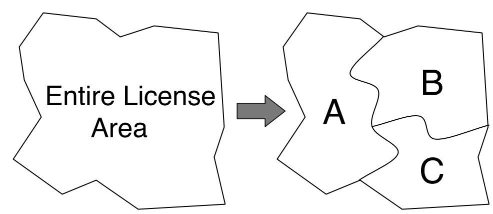

Fig. 1. A license holder partitions its entire license area into three regions, A, B and C. It can sell any of the pieces in the local spectrum market.

a local resource and can only be leased to local users who are within the license region of this channel. In this sense, a *local market* is formed between the license holder and all these local users.

To facilitate such local markets, the FCC has advocated to establish an Internet database to store the information of the currently vacant spectrum as well as its geographic area [2]. An unlicensed user can simply query the database to obtain the list of channels that can be used at the user's location. This database-driven technique enables practical spectrum markets, such as *SpecEx* [3], to take full advantage of the spectrum locality and to provide flexible selling options to attract participation. One of them is *license partitioning* [3], in which a license holder can divide the entire geographic area in which it holds its spectrum license, called its *license area*, into several regions and sell any of the pieces to wireless users, as shown in Fig. 1.

Unfortunately, market locality, as an inherent characteristic of spectrum resources in practice, is seldom mentioned in the literature. Most existing spectrum auctions [4]–[9] are designed in the sense of the *global market* —

- *W. Wang, B. Liang and B. Li are with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON, M5S 3G4. E-mail: weiwang, bli@eecg.toronto.edu, liang@comm.utoronto.ca*
- *Part of this paper has appeared in [1]. This journal version contains substantial revision with new problem models, technical contributions, and simulation results.*

all channels are accessible to all users, no matter where they are. Furthermore, no license partitioning is supported: All channels have to be traded as a whole in the entire license area. These assumptions not only are impractical, but will also seriously degrade the flexibility of selling options, leaving license holders unable or unwilling to join the market. It is typical for channels to be only available to wireless users in some seller-defined geographic regions, rather than the entire license area of a license holder. For example, in Fig. 1, a channel may be vacant in region A, yet utilized by its license holder in region B and C. Without the support of license partitioning (as is the case in global market), a license holder has to decide if it is able to make channels vacant in its entire license area, and channels that are available only in part of the regions are not ready for sale.

To bridge such a gap between the existing literature and practical limitations of geographic spectrum locality, in this paper, we present *District*, a set of new spectrum double auctions that are specifically designed for local spectrum markets. With *District*, a license holder can freely partition its entire license area and either sell or reserve spectrum in local markets, based on their own requirements. Moreover, *District* allows the same channel to be shared by multiple wireless users if no interference occurs.

We believe that it is crucial for *District* to maintain basic properties of economic robustness (truthfulness in particular). As a matter of fact, introducing the notion of local markets imposes non-trivial challenges when economic robustness is to be maintained. Most existing spectrum double auctions [5], [7], [8] are based on McAfee's design [10], which is for the global market only. Their direct extensions, as shown in Sec. 3.6, are either not feasible or leading to fairly inefficient outcomes. To maintain economic robustness, *District* is designed to work effectively in cases with and without *a priori* information about bid distributions. In the former case, *District* extends Myerson's *virtual valuations* [11] to double auctions and designs a market with a discriminatory pricing policy — different auction winners might face different charges or payments. In the latter case, *District* sets a uniform pricing mechanism to charge all winners uniformly. Both mechanisms are proved to be computationally efficient and economically robust. With minor modifications, the design is also applicable to a profit-driven auctioneer aiming to maximize the auction revenue. Extensive simulation studies show that *District* substantially improves spectrum utilization with local markets, and is scalable to large networks.

2 Related Work

Auction serves as an efficient mechanism to price and distribute scarce resources in a market. [12] offers an excellent treatment of general auction theory.

When it comes to spectrum auctions, pioneering works include [13], [14], [15], and [16]. All of them focus on the primary markets, where primary users bid to obtain long-term spectrum rights from the government.

Recently, spectrum auction has received considerable attention in dynamic spectrum access. Some early works in this field include transmit power auctions [17] and spectrum band auctions [18]–[20]. [19] aims to generate maximum revenue for sellers and employ linear programming to model interference constraints. A similar objective is also adopted in [20], where a greedy graph coloring algorithm is used to maximize the revenue in cellular network. [18] also formulates the channel management problem in homogeneous CDMA networks and proposes a spectrum allocation algorithm with linear programming. However, none of these works discusses the strategic behaviours of participants, and the claimed performances are only achieved when all participants truthfully reveal their private information.

The strategic issues of participants are first addressed in [4], in which two truthful spectrum auctions are designed to facilitate as many transactions as possible. [6] further presents a truthful design to maximize the auction revenue. Also, [21] allows sellers to have reservation prices on their idle channels. All these designs are single-sided auctions where only one side, either spectrum buyers or spectrum sellers, has bidding strategies. When the strategic issues extend to both sides, [5] proposes the first spectrum double auction that is truthful. Based on it, [7] designs a double auction for spectrum secondary markets, where participants are all secondary users.

Besides truthfulness, recent works also consider other factors in spectrum auction designs. For example, [8] includes the time domain into strategy considerations and proposes a truthful online auction, while [9] investigates the fairness issue in spectrum allocations.

However, all works above discuss auction designs in the sense of global markets, where all channels to be auctioned off are globally accessible to all spectrum buyers, no matter where they are. Such ignorance of the geographic locality of spectrum resources made them incapable to accommodate the demand of the recent push of database-driven spectrum markets [2], [3], in which channels are traded to local users within the seller-defined license area. A more detailed technical discussion is given in Sec. 3.6.

3 Basic System Model

This section presents the basic system model where a non-profitable auctioneer runs a spectrum double auction for unit-supply sellers and unit-demand buyers. By "unit" we mean each participant either sells or requests a single channel in the local market. In Sec. 6, we further discuss how this model can be extended to scenarios where 1) a buyer can request more than one channel, and 2) the auctioneer is profit-driven aiming to maximize its revenue.

3.1 General Settings

To accommodate spectrum requests that arise dynamically over time, an auctioneer (e.g., the spectrum database administrator) carries out a sequence of double auctions periodically, one in each time slot. Participants, both license holders (sellers) and wireless users (buyers), join the auction asynchronously and submit requests anytime they wish. Yet, their submissions may not be processed immediately until the next auction is conducted, at which time the auction results are computed. Due to the temporal dynamics of spectrum usage, spectrum leases are necessarily ephemeral. A mature business model to facilitate such short-term leases is to adopt an agreed-upon lease cycle, in which a buyer will hold a seller's spectrum license after the transaction is made. Participants may join subsequent auctions to lease channels for additional cycles. Throughout the paper, we focus on auction designs in one time slot, as they are the essential building blocks of the spectrum market and are independent of the value of the time slot or the lease cycle.

In each auction round, every seller has one channel for sale in an indicated license area (a.k.a., local market) in which the channel is vacant (e.g., region A in Fig. 1). Each seller reports the channel, the associated local market (i.e., the available geographic region), and an *ask* to the auctioneer. Each buyer, on the other hand, requests to buy one channel at some geographic location by submitting a *bid* and its location to the auctioneer. All bids and asks are submitted in a sealed manner — no one has access to any information about the others' submissions. After collecting all these submissions, the auctioneer computes the best set of channel transactions to clear the market.

The main challenge is to establish proper payoff schemes and to optimally match buyers and sellers, with the constraints that all channel transactions must be made within local markets, and that no interfering buyers are assigned to the same seller. Fig. 2 illustrates an example of such a double auction with multiple spectrum sellers and buyers in different local markets. Note that the areas of local markets are drawn as circles only for illustration. In fact, they can have arbitrary shapes and may not even be contiguous in general. We finally assume there are M participating license holders and N wireless users.

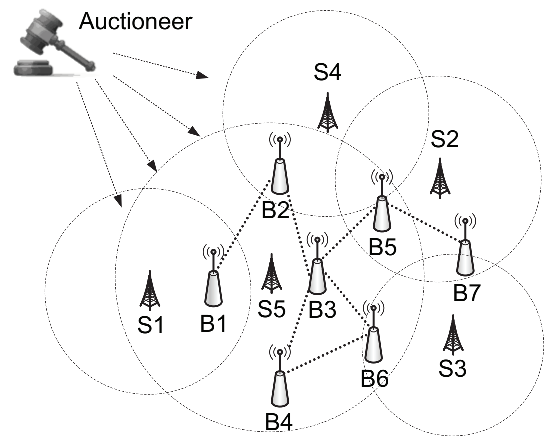

Fig. 2. A spectrum double auction with 7 buyers and 5 sellers. An auctioneer performs the auction among sellers and buyers. Sellers can partition their license areas and sell any pieces of their spectrum in local markets. All license areas for sale are circular in this figure, but can have any shape in general. The dotted lines indicate interference relations among buyers.

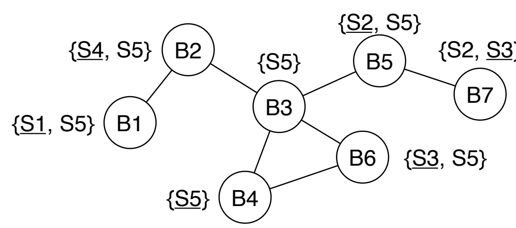

Fig. 3. The graph abstraction $G = (V, E, C)$ of the scenario depicted in Fig. 2. A feasible spectrum assignment is also given. The underlined spectrum is assigned to its corresponding user in the figure.

3.2 Modelling Channel Transactions within Local Markets

Channels should be assigned without introducing interference. We use a conflict graph $G = (V, E)$ to represent interference relations among buyers, where V is the collection of buyers and E is the collection of edges, such that two buyers share an edge if they are in conflict with each other and cannot use the same channel simultaneously. In our example shown in Fig. 2, seven conflicting pairs of buyers are illustrated by dotted lines.

We say seller m and buyer n are *tradable* if n is within m's local market so that it can trade with m. We use C_n to denote the set of tradable sellers of a buyer n. For example, in Fig. 2, $B1$'s tradable sellers are $S1$ and $S5$, i.e., $C_{B1} = \{S1, S5\}$.

Now the network scenario can be represented by a conflict graph G as well as all buyers' tradable sellers $\{C_n\}$. As an example, Fig. 3 illustrates such a representation of the scenario depicted in Fig. 2, with C_{B_i} labeled next to B_i ($i = 1, \dots, 7$).

A channel assignment scheme is *feasible* if all transactions are made between tradable sellers and buyers, and no two buyers sharing an edge in G are assigned to the same seller. A feasible assignment can be equivalently converted to a feasible *graph coloring* scheme by treating tradable sellers C_n as the available colors that can be used to color node n in G. In this sense, a buyer n is assigned to a tradable seller m if node n in G is colored by $m \in C_n$, and vice versa. As an example, one feasible channel assignment of the scenario depicted in Fig. 2 is shown in Fig. 3, with the assigned spectrum underlined in the text. It can be seen that the assignment is exactly a feasible graph coloring scheme. We say a buyer n (a seller m) is a *winner* if node n is colored (color m is used) in G; otherwise, it is a *loser*. For notational convenience, we integrate the available colors of nodes into the conflict graph and denote it by $G = (V, E, C)$, where $C = \{C_n | n \in V\}$.

3.3 Spectrum Double Auction

With the knowledge of G, the auctioneer collects asks (bids) from the sellers (buyers). Denote by a_m and b_n the ask and bid submitted by the seller m and buyer n, respectively. Every seller m has a *true ask* a_m^t, a price that it believes its channel is worth. Every buyer n also has a *true bid* b_n^t, a price quantifying its economic benefit of getting a channel. The value of a_m^t (b_n^t) is the private information of the seller m (buyer n), and is unknown to anyone else, including the auctioneer. Note that the seller m may submit a different ask from its true ask (i.e., $a_m \neq a_m^t$), as long as it believes that this is more beneficial. Similar misreporting strategy may also be adopted by any buyer n (i.e., $b_n \neq b_n^t$).

After collecting all asks $\mathbf{a} = (a_1, \ldots, a_M)$ and bids $\mathbf{b} = (b_1, \ldots, b_N)$, the auctioneer clears the market by computing the channel assignment and winner payoffs. The assignment is represented by a coloring scheme of the conflict graph G, as mentioned before. Every winning seller m is paid p_m by the auctioneer for leasing a channel, while every winning buyer n is charged c_n by the auctioneer. Therefore, the payoffs consist of both the payments $\mathbf{p} = (p_1, \ldots, p_M)$ to sellers and the charges $\mathbf{c} = (c_1, \ldots, c_N)$ to buyers. Then, for each winning pair, the utility of seller m is $u_m^s = p_m - a_m^t$, and that of buyer n is $u_n^b = b_n^t - c_n$. For all losing sellers and buyers, the payments, charges, and corresponding utilities are all zero. The auctioneer gains a revenue, defined as the difference between the total charges and total payments, $\gamma = \sum_n c_n - \sum_m p_m$.

3.4 Economic Requirements

To encourage participation, an auction should satisfy some basic economic requirements [22] as defined below.

Definition 1 (Individual rationality): An auction mechanism is said to be *individually rational* if every participant's utility is nonnegative, i.e., $u_n^b(r_n, b_n) \geq 0$ for all $n = 1, \ldots, N$, and $u_m^s(r_n, b_n) \geq 0$ for all $m = 1, \ldots, M$.

By joining the auction with individual rationality, all participants are guaranteed to be benefited. This property is critical in attracting user participation.

Also, to make the auction self-sustained without any external subsidies, the generated revenue is required to be non-negative.

Definition 2 (Budget balance): An auction mechanism is said to achieve the *ex post* budget balance if its revenue γ is always nonnegative, i.e., $\gamma = \sum_{n=1}^N c_n - \sum_{m=1}^M p_m \geq 0$. Furthermore, if the revenue is nonnegative in expectation, the auction is said to be *ex ante* budget balanced. That is, $\mathbf{E}[\gamma] = \mathbf{E}\left[\sum_{n=1}^N c_n - \sum_{m=1}^M p_m\right] \geq 0$.

Moreover, the designed spectrum double auction is preferred to be truthful.

Definition 3 (Truthfulness): An auction is said to be *truthful* if no participant can expect a higher utility by misreporting its true submission. That is, for every seller m (resp. buyer n), $u_m^s(a_m^t) \geq u_m^s(a_m)$ (resp. $u_n^s(b_n^t) \geq u_n^s(b_n)$) for all a_m (resp. b_n).

As an important economic property, truthfulness brings the following benefits. First, a truthful mechanism helps the auctioneer to gather more accurate market information. We show later in Sec. 5 that utilizing such information enhances the auction performance. Second, truthfulness simplifies all participants' strategies and ensures the basic market fairness. Since no one has the incentive to cheat at its submission, the auction result will not be manipulated by any individual participant. Finally, by the Revelation Principle [11], as long as the revenue is of interest, restricting discussions to truthful auctions does not lose generality.

Following the convention adopted in the existing literature, we say an auction is *economically robust* if it is individually rational, budget balanced (either *ex post* or *ex ante*), and truthful [5], [8].

3.5 Problem Definition

In the basic model, we consider a non-profit auctioneer whose objective is to improve channel utilization by facilitating as many wireless users as possible to access idle channels. For an auction mechanism $\mathcal{M}$, we define the *auction efficiency* as the proportion of winning buyers, i.e., $\eta_{\mathcal{M}} = N_w/N$, where N_w is the number of winning buyers.

With input $G = (V, E, C)$, asks $\mathbf{a}$ and bids $\mathbf{b}$, an auction mechanism $\mathcal{M}$ outputs payments $\mathbf{p}$, charges $\mathbf{c}$, and a colored graph G representing the channel assignment result. Ideally, we would like to find an economically robust auction mechanism that also maximizes the auction efficiency. However, the *impossibility theorem* [23] dictates that the maximal auction efficiency is incompatible with economic robustness. In this work, we view economic robustness as a hard constraint to ensure a well-behaving spectrum market. Hence, we are concerned with the following optimization problem:

$$\begin{aligned} \max_{\mathcal{M}} \quad & \eta_{\mathcal{M}} \\ \text{s.t.} \quad & \mathcal{M} \text{ is economically robust.} \end{aligned} \tag{1}$$

3.6 Challenges of Local Markets

Introducing the notion of local markets imposes non-trivial challenges when economic robustness is to be maintained. Simply extending existing spectrum double auctions, e.g., [5], [7], [8], is either not feasible or leads to fairly inefficient outcomes. In [7] and [8], spectrum reuse is not considered, and a complete conflict graph is assumed. Neither of them is applicable to our system model.

For other auctions designed under a global market that consider spectrum reuse, say, TRUST in [5], it is possible to propose a simple extension for local markets when the traded license areas are of some special shape. For example, suppose all license areas are circular. As shown in Fig. 4, we can partition the entire geographic region into hexagonal cells, each with edge length $R/2$, where R is the minimum radius of all circular license areas. Then, it is always feasible for a buyer to trade with any seller whose license area is centered within the same cell of the buyer. If we limit our discussion to the buyers and sellers within one cell, then TRUST is applicable. Fig. 4 illustrates an example, where the extension is applied to the scenario depicted in

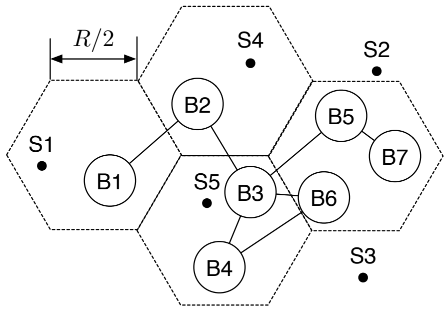

Fig. 4. Applying a simple extension to TRUST [5] to the scenario depicted in Fig. 2.

Fig. 2. It is easy to verify that within each cell all buyers and sellers are tradable to one another. Following this idea, the proposed extension geographically partitions the whole market into a set of submarkets that are independent of each other. It then applies TRUST in every cell to cover the whole market.

However, such an extension is problematic. First, the traded license areas need to be of some special shape, which is usually not the case in practice. Second, buyers are only allowed to trade with sellers within the same cell. As a result, many originally feasible transactions are blocked. For example, in Fig. 4, buyer $B6$ cannot trade with seller $S3$ since they are not covered by the same cell. However, in the original network scenario in Fig. 2, $S3$ and $B6$ are tradable. Similar trade-blocking phenomenon can also be observed among $B5$, $B7$ and $S2$. As a result, $B5$, $B6$ and $B7$ will be left unassigned, independent of their submissions. This intuitive example reveals that the auction efficiency of such simple extension is fairly low, especially when the cell is small. We later verify this problem by our numerical results, shown in Sec. 7.

The failure to simply extend auction mechanisms for global markets requires a new design tailored to local markets. The proposed *District* contains a set of auction mechanisms that are specifically designed to address unique challenges imposed by market locality. We present two alternative designs, *District-D* and *District-U*, for the cases with and without *a priori* knowledge on bid distributions, respectively. An auctioneer without any distribution knowledge can start with *District-U* for a moderate level of auction efficiency, and then switch to *District-D* to pursue a higher level of efficiency after collecting sufficient information about bid distributions.

4 *District-U*: Auction with Uniform Pricing

District-U adopts uniform pricing policies, such that all buyers (sellers) are charged (paid) exactly the same amount of money if they win, without *a priori* knowledge on bids or asks. The basic idea is *trade reduction*: non-profitable trades among low-bid buyers and high-ask sellers are removed, which is critical in maintaining *ex post* budget balance. At a price, however, the auction efficiency is limited due to the reduction of feasible transaction pairs.

4.1 Preliminaries

Without loss of generality, we assume bids and asks are sorted respectively, from the most competitive to the least competitive, i.e., $b_1 \geq b_2 \geq \ldots \geq b_N \geq 0$ and $0 \leq a_1 \leq a_2 \leq \ldots \leq a_M$. We also assume $M \geq N$. Otherwise, we add $N - M$ dummy sellers who ask a_M for their channels but are not tradable to any buyers. Clearly, adding dummy sellers does not affect the auction result.

We introduce the following notations for convenience.

- $G_{m,n}$: A subgraph of $G = (V, E, C)$ composed of only the first m colors (sellers) and the first n nodes (buyers). That is, $G_{m,n} = (V', E', C')$ where $V' = \{1, \ldots, n\}$, $E' = \{e_{ij} \in E | i, j \in V'\}$, and $C' = \{C'_1, \ldots, C'_n\}$ where $C'_j = C_j \cap \{1, \ldots, m\}$.
- *GraphColoring*(G) : A graph coloring algorithm returning a colored G. *District-U* accepts only a *deterministic* graph coloring algorithm with no randomness introduced.

4.2 Mechanism Design

To achieve the trade reduction, *District-U* first decides how many buyers to admit, based on the submitted asks and bids. The auctioneer then computes the set of admitted sellers, and subsequently determines the transaction pairs and corresponding payoffs. The details are given in Algorithm 1.

Algorithm 1 *District-U*

1. $N' = \arg\max_i \{a_{i+1} \leq b_{i+1}\}$
2. $M' = \arg\max_i \{a_i \leq b_{N'+1}\}$
3. *GraphColoring*$(G_{M',N'})$
4. Seller m trades with buyer n if node n is colored by m.
5. **for** each winning buyer n **do**
6. $\quad c_n = b_{N'+1}$
7. **end for**
8. **for** each winning seller m **do**
9. $\quad p_m = b_{N'+1}$
10. **end for**
11. **return** transactions and payoffs ($\mathbf{c}$ and $\mathbf{p}$).

The key to understanding Algorithm 1 is to appreciate its trade reduction nature. As shown in Fig. 5, the algorithm first removes low-bid buyers, i.e., buyers $N' + 1, \ldots, N$, based on the formula of line 1. It then uses the highest bid among the set of removed buyers (i.e., $b_{N'+1}$) as an admission threshold to determine how many sellers to admit: Those who ask higher than the threshold price are all rejected. After the trade reduction, each enrolled buyer bids at least as high as the ask of every admitted seller.

Note that even for the enrolled buyers and sellers, there is no official guarantee that they will eventually win out. The final spectrum assignment will be calculated by running a graph coloring algorithm, and prices are uniformly set for

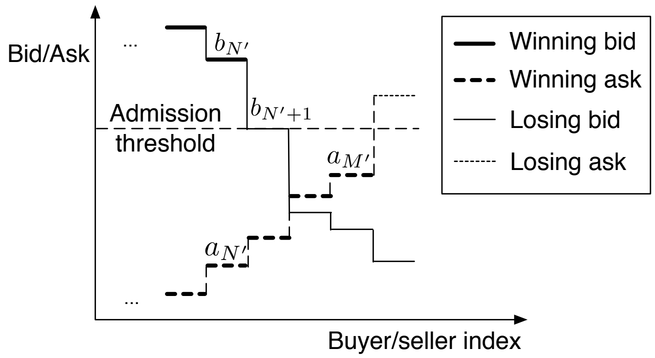

Fig. 5. Illustration of the trade reduction nature in Algorithm 1. The bold lines represent the winning bids and asks. The bids of the enrolled N' buyers are no less than the asks of M' admitted sellers. All others are removed.

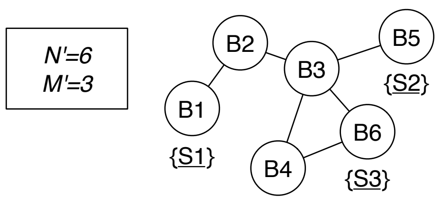

Fig. 6. An example of applying the trade reduction to the scenario depicted in Fig. 2. We assume $N' = 6$ and $M' = 3$. The underlined spectrum is assigned to the corresponding buyer. All winning buyers pay b_{B7} and all winning sellers receive a_{S4}.

both winning buyers and winning sellers. We will see later that this uniform price applied to all buyers (sellers) is the least (highest) submission for them to win. Fig. 6 shows an example of applying Algorithm 1 to the scenario depicted in Fig. 2, where we assume $N' = 6$ and $M' = 3$. We can see that participants $B7, S4, S5$ are directly removed by the auction.

4.3 Economic Properties

Intuitively, one might think that the economic properties of *District-U* should depend on the specific form of $GraphColoring(\cdot)$ adopted by Algorithm 1 in line 3. Surprisingly, the following analysis shows a general result saying that *District-U* is guaranteed to be economically robust, as long as the $GraphColoring(\cdot)$ is a *deterministic* algorithm. That is, with the same input G, algorithm $GraphColoring(\cdot)$ should always produce the same output.

To see this interesting result, we first prove the *ex post* budget balance directly. We then prove that *District-U* is *bid monotonic* with *critical payoffs*, which leads to truthfulness and individual rationality [22], [24].

Proposition 1: District-U is *ex post* budget balanced.

Proof: Denote by x and y the number of winning buyers and sellers, respectively. Since a seller's channel could be spatially reused by multiple buyers, we have $x \geq y$. In this case, the auction revenue is always positive, i.e., $\gamma = \sum_{n=1}^{N} c_n - \sum_{m=1}^{M} p_m = (x - y) \cdot b_{N'+1} \geq 0$. □

Before we proceed to the formal proof of truthfulness and individual rationality, we first introduce two important concepts, the $monotonicity$ and $criticality$. Both of them serve central roles in the proofs.

Definition 4 (Monotonicity): An auction mechanism is *bid monotonic* such that for every buyer n (seller m), if by submitting b_n (a_m) it wins, then by submitting $b'_n > b_n$ ($a'_m < a_m$) it also wins, given the others' submissions remain unchanged.

The interpretation of Definition 4 is quite straightforward: a more competitive submission never hurts a participant's chance to win. An important property associated with a bid monotonic auction is the unique existence of the critical submission defined below.

Definition 5 (Criticality): For *winning* buyer n (seller m), we say b_n^c (a_m^c) is *critical* if n (m) wins by submitting $b_n > b_n^c$ ($a_m < a_m^c$) and loses by submitting $b_n < b_n^c$ ($a_m > a_m^c$), given the others' submissions remain unchanged.

In other words, the critical bid (ask) is the minimum (maximum) submission for a buyer (seller) to win the auction, and is therefore a threshold submission in determining their auction results. Its value depends on other buyers' bids (sellers' asks) and how the winners are selected from the participants.

There is one important result in mechanism design regarding truthfulness and individual rationality, as described in Lemma 1.

Lemma 1 ([22]): A bid monotonic auction is truthful and individually rational if and only if it always charges critical bids from winning buyers and pays critical asks to winning sellers.

Now we show that *District-U* is bid monotonic with critical payoffs, and is hence truthful and individually rational.

Proposition 2: District-U is bid monotonic.

Proof: Suppose buyer n wins by bidding b_n. According to the admission rule stated in line 1, we have $b_n \geq a_n$ and $b_{n+1} \geq a_{n+1}$. Hence, by unilaterally raising the bid up to b'_n, buyer n will be admitted again, as $b'_n > b_n \geq a_n$ and $b_{n+1} \geq a_{n+1}$. We see $G_{M',N'}$ remains unchanged. Since algorithm *GraphColoring* is deterministic, with the same $G_{M',N'}$ as input, the output is always the same, which implies that n wins again.

Similarly, consider a winning seller m decreasing its ask unilaterally from a_m to a'_m. Then $a'_m < a_m \leq b_{N'+1}$. That is, m remains among the bottom M' sellers, and will be admitted again. The rest of the proof is similar to the buyer's case. □

Proposition 3: In *District-U*, $b_{N'+1}$ is the critical submission for all winning buyer n and winning seller m.

Proof: For a winning buyer n, it will remain in the set of the top N' buyers as long as it bids $b_n > b_{N'+1}$, which ensures it to be admitted with the same $G_{M',N'}$. Then n keeps winning due to the unchanged graph coloring result.

On the other hand, suppose n bids $b_n < b_{N'+1}$ and is ranked in the kth place among all buyers, where $N' + 1 \leq$

$k \leq N$. The rankings of bids now become $b_1 \geq \cdots \geq b_{n-1} \geq b_{n+1} \geq \cdots \geq b_k \geq b_n \geq b_{k+1} \geq \cdots \geq b_N$. Since $k+1 > N'+1$, by the definition of N' (i.e., line 1 of Algorithm 1), we have $b_{k+1} < a_{k+1}$. This essentially indicates that buyer n will not be admitted in the new rankings as its next bid b_{k+1} is lower than the corresponding ask a_{k+1} (refer to line 1 of Algorithm 1 for the admission rule).

We now consider the seller's case. For a winning seller m, it will remain in the set of top M' sellers as long as it asks $a_m < b_{N'+1}$, which ensures it to be admitted and win (similar to the buyer's case).

For the case where m asks $a_m > b_{N'+1}$, we see it must be removed because its ask exceeds the admission threshold (see Fig. 5 and line 2 of Algorithm 1). Therefore, m must lose. □

With Propositions 1, 2, 3 and Lemma 1, we conclude the following:

Theorem 1: District-U is economically robust.

It is worth mentioning that the statement of Theorem 1 is independent of the underlying conflict graph. Therefore, a participant will always behave truthfully no matter how its belief on the conflict graph is.

Also note that the adopted graph coloring algorithm plays a key role in *District-U* as it not only affects the auction efficiency, but also dominates the computational complexity. Fortunately, since *District-U* accepts any deterministic coloring algorithms, those computationally efficient designs with good coloring performance could be directly applied. We later evaluate how the choice of the coloring affects the performance of *District-U* in Sec. 7.

In conclusion, *District-U* is designed as a suitable starting mechanism for the auctioneer without *a priori* information on bids or asks.

5 *District-D*: Auction with Discriminatory Pricing

When bid and ask distributions are available, one can expect higher efficiency via *District-D*, an auction with discriminatory pricing policies (i.e., winners have different payoffs). We begin by using Myerson's *virtual valuations* [11] to characterize the expected revenue of a truthful spectrum auction $\mathcal{M}$. We show that an economically robust $\mathcal{M}$ is equivalent to a graph coloring with a weighted sum non-negative in expectation. We design *District-D* based on this result and show that the design is computationally efficient.

5.1 Preliminaries

Assume all buyers (sellers) bid (ask) independently, but possibly under different distributions. For every buyer n, denote its bid distribution function as $F_n^b(b_n)$ and the corresponding density function as $f_n^b(b_n)$. We similarly define $F_m^s(a_m)$ and $f_m^s(a_m)$ for every seller m. Since $\mathcal{M}$ is designed to be truthful, we do not discriminate the true bid (ask) and the submitted bid (ask) in the following context.

In [11], Myerson defines *virtual valuations* for buyers in a single-sided auction. We extend this idea to double auctions and apply it to our design. Formally, we define

$$\psi_m(a_m) = a_m + \frac{F_m^s(a_m)}{f_m^s(a_m)} \tag{2}$$

as the virtual valuation of seller m with ask a_m, and

$$\phi_n(b_n) = b_n - \frac{1 - F_n^b(b_n)}{f_n^b(b_n)} \tag{3}$$

as the virtual valuation of buyer n with bid b_n. We assume *regular* distributions [11], i.e., all $\phi_n(\cdot)$ and $\psi_m(\cdot)$ are increasing functions[1].

Let $\mathbf{v} = (a_1, \ldots, a_M, b_1, \ldots, b_N)$ be the vector of submissions. Denote by $\gamma_{\mathcal{M}}(\mathbf{v}, G)$ the revenue of auction $\mathcal{M}$ with submissions $\mathbf{v}$ and the conflict graph G as the input. When there is no confusion, we simply write $\gamma_{\mathcal{M}}(\mathbf{v}, G)$ as $\gamma(\mathbf{v})$. The following lemma shows that the expected revenue can be fully characterized by the virtual valuations of all winners. The proof is similar to [11] and is given in the appendix[2].

Lemma 2: Given a truthful $\mathcal{M}$ and G, let $x_n(\mathbf{v}) = 1$ if n wins, i.e., n is colored, and $x_n(\mathbf{v}) = 0$ otherwise. Let $y_m(\mathbf{v})$ be similarly defined for a seller m. Then

$$\mathbf{E}_{\mathbf{v}}\left[\gamma(\mathbf{v})\right] = \mathbf{E}_{\mathbf{v}}\left[\sum_{n=1}^{N} \phi_n(b_n) x_n(\mathbf{v}) - \sum_{m=1}^{M} \psi_m(a_m) y_m(\mathbf{v})\right] . \tag{4}$$

Lemma 2 reveals an important fact — dealing with virtual valuations is equivalent to dealing with submitted bids (asks), in terms of the expected revenue.

Introducing $\phi_n(\cdot)$ and $\psi_m(\cdot)$ greatly simplifies the auction design problem. Suppose the conflict graph G is given. For a buyer n with bid b_n, we assign $\phi_n(b_n)$ as the node weight to node n. For seller m with ask a_m, we assign $\psi_m(a_m)$ as the color weight to color m. We define the *weight of a colored graph* G as

$$W(G) = \sum_{n=1}^{N} \phi_n(b_n) \cdot x_n - \sum_{m=1}^{M} \psi_m(a_m) \cdot y_m . \tag{5}$$

Here $x_n = 1$ if n is colored, and $x_n = 0$ otherwise; $y_m = 1$ if m is used to color, and $y_m = 0$ otherwise. By Lemma 2, we have

$$\mathbf{E}_{\mathbf{v}}\left[\gamma_{\mathcal{M}}(\mathbf{v}, G)\right] = \mathbf{E}_{\mathbf{v}}\left[W(G)\right] . \tag{6}$$

Therefore, achieving the *ex ante* budget balance is equivalent to maintaining the weighted sum non-negative in expectation, i.e., $\mathbf{E}[W(G)] \geq 0$.

Recall that the ultimate goal of our design is to pursue a high auction efficiency. We can now rewrite the objective problem (1) as follows:

$$\begin{aligned} \max_{\mathcal{M}} \quad & \mathbf{E}\left[\eta_{\mathcal{M}}\right] \\ \text{s.t.} \quad & \mathbf{E}[W(G)] \geq 0 , \\ & \mathcal{M} \text{ is truthful and individually rational} . \end{aligned} \tag{7}$$

1. This is not a restrictive assumption, as many important distributions satisfy the regularity assumption [11], including uniform, exponential, normal, etc.

2. The appendix is given in a supplementary document as per the TMC submission guidelines.

We see from the analysis above that the expected revenue of a truthful mechanism can be fully characterized by winners only, independent of their payoffs. In other words, for truthful spectrum auctions, the winner-determination algorithm (i.e., graph coloring) characterizes the pricing scheme, and serves a key role in *District-D*'s design.

5.2 Winner Determination

Based on our model, calculating the winning buyers and sellers is, in essence, to calculate a graph coloring scheme. From (7), we see that such a coloring scheme should color as many nodes as possible, while keeping the weighted sum non-negative in expectation. Due to the intractability of this problem, our design is heuristic-based.

In the proposed heuristic algorithm, at every iteration, we pick a feasible buyer-seller pair with the maximum marginal revenue measured by the virtual valuation. We then add the pair's marginal revenue to the accumulated revenue. If a deficit (i.e., the resulted revenue is negative) occurs, the pair is rejected and the algorithm terminates. Otherwise, the pair is accepted, and we proceed to the next iteration and repeat the above procedure.

For convenience, we introduce the following notations before presenting the formal algorithm in Algorithm 2.

- $\mathbf{S}$ – Set of all sellers $\mathbf{S} = \{1, \ldots, M\}$.
- $\mathbf{B}$ – Set of all buyers $\mathbf{B} = \{1, \ldots, N\}$.
- $\mathcal{T}$ – Round-by-round record of transactions already made by the winner-determination algorithm.
- $\mathcal{T}^b$ – Set of winning buyers corresponding to $\mathcal{T}$.
- $\mathcal{T}^s$ – Set of winning sellers corresponding to $\mathcal{T}$.
- $\Delta_{m,n}(\mathcal{T}, a_m, b_n)$ – Marginal revenue generated by assigning m to n, given $\mathcal{T}$, seller m's ask a_m, and buyer n's bid b_n. If the assignment is feasible, then by (4), we have
$$\Delta_{m,n}(\mathcal{T}, a_m, b_n) = \phi_n(b_n) - \psi_m(a_m) I_{m \notin \mathcal{T}^s} , \quad (8)$$
where $I_\alpha = 1$ if α is true, $I_\alpha = 0$ otherwise. When the assignment is not feasible, we simply define $\Delta_{m,n}(\mathcal{T}, a_m, b_n) = -\infty$.
- *MaxMarginalRev*($\mathcal{T}$) – Given $\mathcal{T}$, calculate the transaction (m, n) with the maximum marginal revenue among all feasible transactions, i.e., $\Delta_{m,n}(\mathcal{T}, a_m, b_n) = \max_{i \in \mathbf{S}, j \in \mathbf{B}} \Delta_{i,j}(\mathcal{T}, a_i, b_j)$.

Algorithm 2 *District-D* Winner Determination

```
Initialization: γ ← 0, T ← ∅, and stop ← false.
while stop = false do
    Δ_{m,n} ← MaxMarginalRev(T)
    if γ + Δ_{m,n} ≥ 0 then
        γ ← γ + Δ_{m,n}
        Add (m, n) to T, color G accordingly.
    else
        stop ← true
    end if
end while
return T
```

To better understand Algorithm 2, one can refer to Fig. 7 for an example.

Our winner-determination algorithm guarantees the budget balance, as stated by the following proposition.

Proposition 4: *District-D* is *ex ante* budget balanced.

Proof: Given conflict graph G, for all asks $\mathbf{a}$ and all bids $\mathbf{b}$, Algorithm 2 colors G with a non-negative weighted sum, i.e., $\gamma = W(G) \geq 0$. Hence $\mathbf{E}[W(G)] = \mathbf{E}[\gamma_{\mathcal{M}}(\mathbf{v}, G)] \geq 0$, where the equality holds because of (6). Since this statement holds for every G, we conclude the proof. □

Also, our design has the potential to achieve truthfulness and individual rationality. The key to this point is to show that the design is bid monotonic. Once this is the case, the pricing scheme will be an algorithm to calculate the critical submissions for all winners.

Proposition 5: Algorithm 2 is bid monotonic.

The formal proof of Proposition 5 is presented in the appendix. Intuitively, if buyer n wins and trades with seller m by bidding b_n, then raising the bid to $b'_n > b_n$ increases the marginal revenue generated by transaction (m, n), i.e., $\Delta_{m,n}(\mathcal{T}, a_m, b'_n) \geq \Delta_{m,n}(\mathcal{T}, a_m, b_n)$, as can be verified by referring to (8). Since (m, n) is already the most profitable assignment when n bids b_n, we see that in the new submission, (m, n) remains the most profitable choice for the auctioneer and will be selected again. Therefore, n wins by bidding higher than b_n. Similar argument also applies to winning seller m.

Now by Lemma 1, to achieve truthfulness and individual rationality, the pricing mechanism of *District-D* needs to calculate the critical submissions in a computationally efficient manner. We present the design in the next sections.

5.3 Buyer Pricing

By Definition 5, we see that the critical bid is the minimum submission required to win. The basic rationale is that to win the auction, there is no need to bid as high as possible. Instead, it suffices to win if one can do better than its competitors. Following this idea, we first remove buyer n from bidding. We then conduct winner determination to obtain the winners list and see the winning competitors' bids. Buyer n can win as long as its bid is higher than the one submitted by the weakest competitor. The detailed procedure is given in Algorithm 3. It is worth mentioning that if the equation in line 4 (resp. line 8) has no solution, then the value of $b^l_{(i)}$ (resp. $b^{k+1}_{(i)}$) is set to be ∞.

Let $\mathcal{T}$ be the transaction list generated by the winner-determination algorithm (Algorithm 2). Let $\mathcal{T}_l$ be the first l transactions in $\mathcal{T}$, i.e., $\mathcal{T}_l = \{(i_1, j_1), \ldots, (i_l, j_l)\}$, where (i_l, j_l) is the lth transaction made by Algorithm 2 between the seller i_l and the buyer j_l. Denote by $\mathcal{T}^s_l$ the winning sellers associated with $\mathcal{T}_l$.

Proposition 6: For every winning buyer n, Algorithm 3 returns its critical bid c_n.

Proof: We first prove that n wins by bidding higher than c_n, i.e., $b_n > c_n$. It suffices to consider two cases:

Case 1: c_n is finalized in the first k loops, i.e., $c_n = b^l = b^l_{(i)}$ for some $l \leq k$ and $i \in \mathbf{S}$. For n bidding

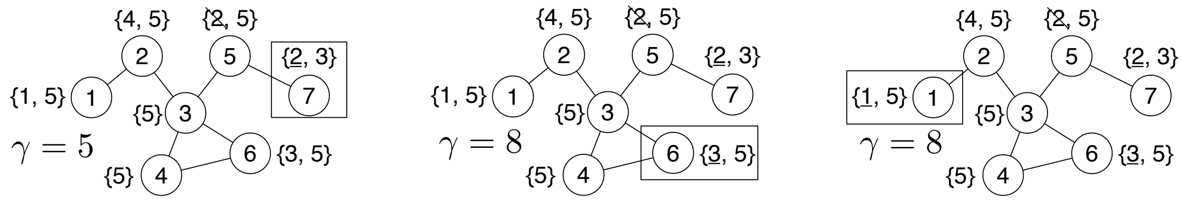

(a) The first iteration. $\mathcal{T} = \{(2,7)\}$.

(b) The second iteration. $\mathcal{T} = \{(2,7),(3,6)\}$.

(c) The third iteration. $\mathcal{T} = \{(2,7),(3,6),(1,1)\}$.

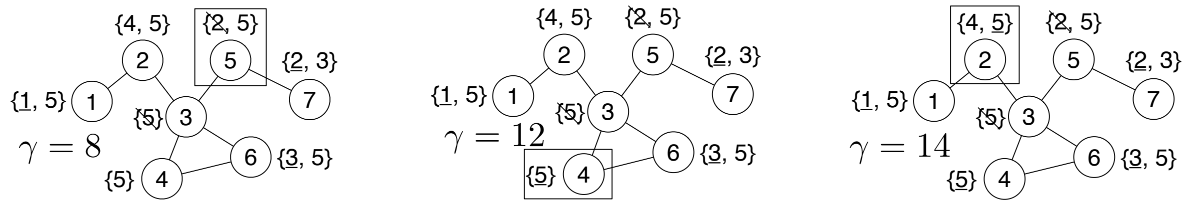

(d) The fourth iteration. $\mathcal{T} = \{(2,7),(3,6),(1,1),(5,5)\}$.

(e) The fifth iteration. $\mathcal{T} = \{(2,7),(3,6),(1,1),(5,5),(5,4)\}$.

(f) The final iteration. $\mathcal{T} = \{(2,7),(3,6),(1,1),(5,5),(5,4),(5,2)\}$.

Fig. 7. An example of Algorithm 2. The network scenario is the same as that depicted in Fig. 2. For simplicity, the virtual valuations are set as follows: $\psi_m = m$ and $\phi_n = n$. We replace labels for buyers and sellers by their virtual valuations in figures. (a) to (f) illustrate the iteration process of Algorithm 2.

Algorithm 3 *District-D* Pricing for a Winning Buyer n

1. Remove n and run Algorithm 2 to obtain the transaction (*seller, buyer*) list $\mathcal{T} = \{(i_1, j_1), \ldots, (i_k, j_k)\}$.
2. $c_n \leftarrow \infty$, and $\gamma \leftarrow 0$
3. **for** $l = 1$ to k **do**
4. $\quad b^l \leftarrow \min_{i \in \mathbf{S}} b^l_{(i)}$, where $b^l_{(i)}$ solves the following equation: $\Delta_{i,n}(\mathcal{T}_{l-1}, a_i, b^l_{(i)}) = \Delta_{i_l,j_l}(\mathcal{T}_{l-1}, a_{i_l}, b_{j_l})$
5. $\quad c_n \leftarrow \min\{c_n, b^l\}$
6. $\quad \gamma \leftarrow \gamma + \Delta_{i_l,j_l}(\mathcal{T}_{l-1}, a_{i_l}, b_{j_l})$
7. **end for**
8. $b^{k+1} \leftarrow \min_{i \in \mathbf{S}} b^{k+1}_{(i)}$, where $b^{k+1}_{(i)}$ solves the following equation: $\Delta_{i,n}(\mathcal{T}_k, a_i, b^{k+1}_{(i)}) = -\gamma$
9. $c_n \leftarrow \min\{c_n, b^{k+1}\}$
10. **return** c_n

$b_n > c_n = b^l_{(i)}$, the worst case is that it loses in the first $l-1$ rounds. But in the lth round, $\Delta_{i,n}(\mathcal{T}_{l-1}, a_i, b_n) > \Delta_{i,n}(\mathcal{T}_{l-1}, a_i, b^l_{(i)}) = \Delta_{i_l,j_l}(\mathcal{T}_{l-1}, a_{i_l}, b_{j_l})$, where the first inequality holds due to the increasing $\phi_n(\cdot)$ and (8). This implies that making the transaction (i, n) would generate more marginal revenue than selecting (i_l, j_l). Since (i_l, j_l) is already the most profitable transaction when n is absent, we conclude that (i, n) maximizes the marginal revenue when n joins the auction. Therefore, n wins by being selected to trade with seller i.

Case 2: $c_n = b^{k+1} = b^{k+1}_{(i)}$ for some $i \in \mathbf{S}$. For n bidding $b_n > c_n = b^{k+1}_{(i)}$, the worst case is that it loses in the first k rounds. However, n can still trade with seller i after the first k rounds, with the marginal revenue $\Delta_{i,n}(\mathcal{T}_k, a_i, b_n) > \Delta_{i,n}(\mathcal{T}_k, a_i, b^{k+1}_{(i)}) = -\gamma$. Therefore, adding the transaction pair (i, n) to the auction results makes the total revenue remain positive, i.e., $\Delta_{i,n}(\mathcal{T}_k, a_i, b_n) + \gamma > 0$. Based on the winner determination algorithm (Algorithm 2), n wins and trades with seller i.

Next, if n bids less than c_n (i.e., $b_n < c_n$), then it loses in the first k rounds. For any $l = 1, \ldots, k$ and any seller $i \in \mathbf{S}$, we have

$$\begin{aligned} \Delta_{i,n}(\mathcal{T}_{l-1}, a_i, b_n) &< \Delta_{i,n}(\mathcal{T}_{l-1}, a_i, b^l) \\ &\le \Delta_{i,n}(\mathcal{T}_{l-1}, a_i, b^l_{(i)}) \\ &= \Delta_{i_l,j_l}(\mathcal{T}_{l-1}, a_{i_l}, b_{j_l}) \ . \end{aligned}$$

Here the equality holds because of line 4. This essentially indicates that buyer n loses in the first k rounds as any transactions involving it does not generate the optimal marginal revenue.

Moreover, even if $b^{k+1} < \infty$ after k rounds, n loses and cannot trade with any seller $i \in \mathbf{S}$. Otherwise, the total revenue would become negative: $\Delta_{i,n}(\mathcal{T}_k, a_i, b_n) + \gamma \le \Delta_{i,n}(\mathcal{T}_k, a_i, b^{k+1}) + \gamma \le \Delta_{i,n}(\mathcal{T}_k, a_i, b^{k+1}_{(i)}) + \gamma = 0$. $\square$

5.4 Seller Pricing

For sellers, the analysis on pricing is similar to buyers'. Seller m only needs to ask for lower than its competitors to win the auction. We first remove m and run the winner determination to see its winning competitors' asks. Seller m can win by asking for lower than these competitors. The detailed design is given in Algorithm 4. Note that if the equation in line 4 (resp. line 8) has no solution, then the value of $a^l_{(j)}$ (resp. $a^{k+1}_{(j)}$) is set to be $-\infty$.

Proposition 7: For every winning seller m, Algorithm 4 returns its critical ask p_m.

The proof of Proposition 7 is similar to the proof of Proposition 6 and is given in the appendix.

Algorithm 4 *District-D* Pricing for a Winning Seller m

1. Remove m and run Algorithm 2 to obtain the transaction (*seller, buyer*) list $\mathcal{T} = \{(i_1, j_1), \ldots, (i_k, j_k)\}$.
2. $p_m \leftarrow -\infty$, and $\gamma \leftarrow 0$
3. **for** $l = 1$ to k **do**
4. $\quad a^l \leftarrow \max_{j \in \mathbf{B}} a^l_{(j)}$, where $a^l_{(j)}$ solves the following equation: $\Delta_{m,j}(\mathcal{T}_{l-1}, a^l_{(j)}, b_j) = \Delta_{i_l, j_l}(\mathcal{T}_{l-1}, a_{i_l}, b_{j_l})$
5. $\quad p_m \leftarrow \max\{p_m, a^l\}$
6. $\quad \gamma \leftarrow \gamma + \Delta_{i_l, j_l}(\mathcal{T}_{l-1}, a_{i_l}, b_{j_l})$
7. **end for**
8. $a^{k+1} \leftarrow \max_{j \in \mathbf{B}} a^{k+1}_{(j)}$, where $a^{k+1}_{(j)}$ solves the following equation: $\Delta_{m,j}(\mathcal{T}_k, a^{k+1}_{(j)}, b_j) = -\gamma$
9. $p_m \leftarrow \max\{p_m, a^{k+1}\}$
10. **return** p_m

5.5 Economic Properties and Computational Efficiency

From Proposition 5, 6 and 7, we see that *District-D* is bid monotonic and generates critical submissions. Therefore, by Lemma 1, it is truthful and individually rational. Noting that *District-D* is also *ex ante* budget balanced by Proposition 4, we conclude that:

Theorem 2: District-D is *economically robust*.

We now analyze the time complexity of *District-D*. In Algorithm 2, one transaction is made in each round of the loop, and the loop runs at most N rounds for N transactions. The complexity of each round is dominated by *MaxMarginalRev*$(\cdot)$, which takes $O(M^2N)$ time to finalize[3]. We hence need $O(M^2N^2)$ time for Algorithm 2. Note that Algorithm 2 also dominates the complexity of Algorithm 3 and 4, where the former runs N times while the later runs M times to calculate prices for all winners. We conclude that *District-D* completes within $O(M^2N^3 + M^3N^2)$ time.

6 Extensions

For now, all our discussions are based on a basic model where each buyer requests only a single channel and the auctioneer is non-profit — its ultimate goal is to maximize the auction efficiency, not the revenue. In this section, we show that with minor modifications, previously developed techniques also apply to more general cases, in which 1) buyers can bid for multiple channels, and 2) the auctioneer aims to maximize the auction revenue.

Note that a more general model should also include multi-supply sellers selling different channels idle in different geographic regions, at different prices. Each seller jointly optimizes the ask prices of different channels to maximize the overall utility. Since every seller has high-dimensional strategies, the problem is essentially a combinatorial mechanism design problem [22]. We are not aware of any revenue characterization result reported in this setting. Without this result, it would be very difficult to ensure budget balance, as there is no way to characterize the auction revenue. For these reasons, we leave this general scenario as an open problem for future investigation.

3. A simple implementation is to go through all tradable transactions without conflicting with previously made trades, and select the one with the maximum marginal revenue. There are at most MN such transactions, each requiring at most M comparisons to clear the conflict relation. We hence need $O(M^2N)$ time for *MaxMarginalRev*$(\cdot)$.

6.1 Multi-Demand Buyers

6.1.1 Models

Suppose every buyer n requests r_n channels at its location and has a true bid b^t_n for accessing each one. Buyer n reports r_n and submits a bid b_n to maximize its utility. No *partial fulfillment* is accepted: A buyer is either rejected or has all requested r_n channels being allocated. We define buyer n's utility $u^b_n(b_n)$ as follows:

$$u^b_n(b_n) = \begin{cases} r_n(b^t_n - c_n)\,, & \text{if } n \text{ wins;} \\ 0\,, & \text{otherwise,} \end{cases} \tag{9}$$

where c_n is the charged per-channel price. A buyer's problem is to find an optimal submission to maximize its utility, i.e., $\max_{b_n} u^b_n(b_n)$. We note that this multi-demand buyer model has also been adopted in [4], [21]. The seller model is the same as in Sec. 3.

The auctioneer's problem is to design a spectrum double auction to maximize the auction efficiency $\eta_{\mathcal{M}}$ while maintaining economic robustness. Here, $\eta_{\mathcal{M}}$ is defined as the ratio of the fulfilled channel requests to the total requests, i.e., $\eta_{\mathcal{M}} = \sum_{n \in W_b} r_n / \sum_n r_n$, with W_b being the set of all winning buyers.

Note that the auctioneer has to ensure a feasible channel assignment such that no two interfering buyers shares the same channel. To this end, we first extend the conflict graph $G(V, E, C)$ defined in Sec. 3 as follows. For buyer n requesting r_n channels, we treat it as r_n virtual buyers each requesting a different channel. We therefore split node n in G to r_n subnodes, each having exactly the same available colors C_n and interfering relations as its supernode n has. In particular, if buyer n interferes with n', then every subnode of n has an edge connected to every subnode of n'. Besides, to ensure that all r_n channels allocated to buyer n are different, we add an interfering edge between every two subnodes of it. For convenience, denote by $\bar{G}$ the extended conflict graph of the original G. A feasible channel assignment therefore corresponds to a coloring scheme on an extended conflict graph $\bar{G}$.

6.1.2 Extending District-U

In Algorithm 1, after the trade reduction, instead of coloring the original conflict graph $G_{M',N'}$ (see line 3), we color the extended graph $\bar{G}_{M',N'}$, and buyer n wins if all its r_n subnodes are colored. The pricing scheme remains the same as the original one: Every winning seller is paid $b_{N'+1}$ and every winning buyer is charged $b_{N'+1}$ for each traded channel.

It is easy to verify that both the claims and proofs of Proposition 1, 2, and 3 trivially apply to the above extensions, leading to the properties of truthfulness, bid monotonicity, and critical submissions. Therefore, we have the following theorem.

Theorem 3: The extended *District-U* is economically robust.

6.1.3 Extending District-D

Since a winning buyer n now trades with r_n sellers, by Lemma 2, in expectation, accommodating its channel requests generates $r_n\phi_n(b_n)$ revenue for the auctioneer. In fact, we see that Lemma 2 can be trivially extended to characterize the expected auction revenue as

$$\mathbf{E_v}\left[\gamma(\mathbf{v})\right] = \mathbf{E_v}\left[\sum_{n=1}^{N} r_n\phi_n(b_n)x_n(\mathbf{v}) - \sum_{m=1}^{M}\psi_m(a_m)y_m(\mathbf{v})\right]. \tag{10}$$

Therefore, being budget balanced is equivalent to maintaining the RHS of (10) nonnegative.

Some notations are introduced before presenting the formal designs. Denote by (S, n) the transactions made between buyer n and a set of sellers $S = \{s_1, \ldots, s_{r_n}\}$. Let $\mathbf{S}$, $\mathbf{B}$, $\mathcal{T}$, $\mathcal{T}^b$, $\mathcal{T}^s$, $\mathcal{T}_l$, and $\mathcal{T}_l^s$ be similarly defined as they are in Sec. 5.2 and 5.3. Denote by $\Delta_{S,n}(\mathcal{T})$ the marginal revenue generated by adding transactions (S, n) to the existing transaction list $\mathcal{T}$. By (10), when adding (S, n) to $\mathcal{T}$ is feasible, we have

$$\Delta_{S,n}(\mathcal{T}) = r_n\phi_n(b_n) - \sum_{i=1}^{r_n}\psi_{s_i}(a_{s_i})I_{s_i\notin\mathcal{T}^s}. \tag{11}$$

Otherwise, $\Delta_{S,n}(\mathcal{T}) = -\infty$. Sometimes, $\Delta_{S,n}(\mathcal{T})$ is also written as $\Delta_{S,n}(\mathcal{T}, a_m)$ ($\Delta_{S,n}(\mathcal{T}, b_n)$) if seller m's ask (buyer n's bid) is of interest. Finally, let *MaxMarginalRev*$(\mathcal{T})$ return the feasible transactions (S, n) with the maximum marginal revenue, given the already decided transaction list $\mathcal{T}$, i.e., *MaxMarginalRev*$(\mathcal{T}) = \max_{n\in\mathbf{B}, S\subset\mathbf{S}, |S|=r_n}\Delta_{S,n}(\mathcal{T})$.

Winner determination. The extended winner determination algorithm follows the same logic flow of Algorithm 2, with $\Delta_{m,n}$ and (m, n) being replaced by $\Delta_{S,n}$ and (S, n), respectively. Intuitively, the auctioneer greedily adds transactions (S, n) with the maximum marginal revenue to its output, until the accumulated revenue is no longer positive, at which time it stops.

Buyer pricing. With the same design idea of Algorithm 3, we extend the buyer pricing to Algorithm 5. Note that in line 4 (resp. line 8), if the equation has no solution, we simply set $b^l_{(S)} = \infty$ (resp. $b^{k+1}_{(S)} = \infty$).

Seller pricing. Algorithm 6 presents the design of the extended seller pricing. If the equation in line 4 (resp. line 8) has no solution, we simply set $a^l_{(j)} = -\infty$ (resp. $a^{k+1}_{(j)} = -\infty$).

Following the same analyses of Proposition 4 and 5, we can easily show that the winner determination is *ex ante* budget balanced and bid monotonic, while both the buyer and seller pricing return critical submissions. Therefore, we have the following theorem. The proofs are given in the appendix.

Theorem 4: The extended *District-D* is economically robust.

Algorithm 5 Extended *District-D* Pricing for a Winning Buyer n

1. Remove n and run Algorithm 2 to obtain the transaction list $\mathcal{T} = \{(S_1, j_1), \ldots, (S_k, j_k)\}$.
2. $c_n \leftarrow \infty$, and $\gamma \leftarrow 0$
3. **for** $l = 1$ to k **do**
4. $\quad b^l \leftarrow \min_{S\subset\mathbf{S}, |S|=r_n} b^l_{(S)}$, where $b^l_{(S)}$ solves the equation $\Delta_{S,n}(\mathcal{T}_{l-1}, b^l_{(S)}) = \Delta_{S_l,j_l}(\mathcal{T}_{l-1}, b_{j_l})$.
5. $\quad c_n \leftarrow \min\{c_n, b^l\}$
6. $\quad \gamma \leftarrow \gamma + \Delta_{S_l,j_l}(\mathcal{T}_{l-1}, b_{j_l})$
7. **end for**
8. $b^{k+1} \leftarrow \min_{S\subset\mathbf{S}, |S|=r_n} b^{k+1}_{(S)}$, where $b^{k+1}_{(S)}$ solves the equation $\Delta_{S,n}(\mathcal{T}_k, b^{k+1}_{(S)}) = -\gamma$
9. $c_n \leftarrow \min\{c_n, b^{k+1}\}$
10. **return** c_n

Algorithm 6 Extended *District-D* Pricing for a Winning Seller m

1. Remove m and run Algorithm 2 to obtain the transaction list $\mathcal{T} = \{(S_1, j_1), \ldots, (S_k, j_k)\}$.
2. $p_m \leftarrow -\infty$, and $\gamma \leftarrow 0$
3. **for** $l = 1$ to k **do**
4. $\quad a^l \leftarrow \max_{j\in\mathbf{B}} a^l_{(j)}$, where $a^l_{(j)}$ solves the equation $\max_{S\subset\mathbf{S}, |S|=r_j, m\in S}\{\Delta_{S,j}(\mathcal{T}_{l-1}, a^l_{(j)})\} = \Delta_{S_l,j_l}(\mathcal{T}_{l-1})$
5. $\quad p_m \leftarrow \max\{p_m, a^l\}$
6. $\quad \gamma \leftarrow \gamma + \Delta_{S_l,j_l}(\mathcal{T}_{l-1})$
7. **end for**
8. $a^{k+1} \leftarrow \max_{j\in\mathbf{B}} a^{k+1}_{(j)}$, where $a^{k+1}_{(j)}$ solves the equation $\max_{S\subset\mathbf{S}, |S|=r_j, m\in S}\{\Delta_{S,j}(\mathcal{T}_k, a^{k+1}_{(j)})\} = -\gamma$
9. $p_m \leftarrow \max\{p_m, a^{k+1}\}$
10. **return** p_m

Note that all extensions above remain computationally efficient, taking $O(M^2N^3 + M^3N^2)$ time to complete. The detailed analysis is also given in the appendix.

6.2 Profit-Driven Auctioneer

Previous discussions assume a non-profit auctioneer aiming to maximize the auction efficiency. Such an unselfish auctioneer could be governmental or public agencies such as the FCC in US. However, when the auction is operated by some profit-driven company, such as *Spectrum Bridge, Inc.*, then the auctioneer's objective is to maximize its own revenue. In this case, the auctioneer runs *District-U* first. After accumulating sufficient knowledge of bid distributions, it switches to *District-D* with the following minor modifications.

By Lemma 2, maximizing the expected auction revenue is equivalent to maximizing the RHS of (4), which is apparently an NP-Complete problem. We hence adopt a heuristic solution with greedy designs: Sequentially add a transaction (m, n) with the maximum marginal revenue $\Delta_{m,n}$ to the auction output, until there is no profitable transaction to make (i.e., $\Delta_{m,n} < 0$). In particular, we change the "if" condition from $\gamma + \Delta_{m,n} \geq 0$ (see line 4 of Algorithm 2) to $\Delta_{m,n} \geq 0$. Some adjustments are also needed in the pricing design: We replace $-\gamma$ in line 8 of Algorithm 3 with 0. The same replacement is also taken in the same line of Algorithm 4.

With similar proofs to Proposition 5, 6 and 7, one can easily show that the extension above is bid monotonic and charges (pays) the winners their critical submissions. Therefore, we have the following conclusion.

Theorem 5: The extension of *District-D* for a profit-driven auctioneer is truthful and individually rational.

Finally, for multi-demand buyers, the extensions presented in Sec. 6.1 are also applicable, and Theorem 5 extends to this scenario as well. The complete design is given in the appendix.

7 Simulation Results

We evaluate the performance of *District* with extensive simulations. Buyers are uniformly distributed in a 1×1 geographical region. Two buyers interfere with each other if their Euclidean distance is less than 0.1. Every seller indicates a license area to sell. The license area is assumed to be circular, with radius uniformly distributed between 0.2 and 0.5 and center uniformly located in the entire region. All bids and asks are normalized and follow the uniform distribution in the range of $[0, 1]$. Since *District* is truthful by design, these are also true bids and asks.

We evaluate the auction efficiency as the main performance metric of *District*, as our design is already proven to be economically robust. We also present numerical results for the generated revenue to study the relation between the auction efficiency and the revenue. All results presented in this section are based on the basic scenario where every participant trades only one channel in the local market. For the extended case where a buyer asks for multiple channels, we observe similar performance trends, which is reasonable since the extensions adopt similar design ideas to the original ones. Each result obtained below has been averaged over 10000 runs.

7.1 District-U

We first investigate how the performance of *District-U* is affected by the choice of the deterministic graph coloring algorithm *GraphColoring*$(\cdot)$. We compare three sequential coloring algorithms with different coloring orders, i.e., the fixed order where nodes with the least node ID is colored first, the least uncolored neighbors where a node with the least uncolored neighbors is colored first, and Brélaz's DSATUR where a node with the least available colors is colored first [25]. Fig. 8a illustrates the comparison results of their performances. Interestingly, we see that all three coloring algorithms achieve similar auction efficiency. This suggests that the performance of *District-U* is less dependent on the specific design of *GraphColoring*$(\cdot)$. However, adopting a good coloring algorithm, such as DSATUR, does benefit the auction efficiency.

Since DSATUR slightly outperforms the other two coloring, our following simulations are carried out based on *District-U* with DSATUR. Fig. 8b and 8c further illustrate its mean and standard deviation of the auction efficiency, respectively, while Fig. 8d presents the mean revenue.

We see that the mean auction efficiencies in all four experiments exhibit similar trend in their performance curves: Higher auction efficiency is achieved as more sellers join the auction, while the marginal enhancement is decreasing. On the other hand, when the channel supply is maintained unchanged, adding more buyers to the market degrades the auction efficiency, as the fixed amount of channel supply is insufficient to accommodate the increasing demand. When the channel supply and demand are of the same level (i.e., the number of buyers is equal to the number of sellers), the auction efficiency is around 50%.

It is worth mentioning that the maximum revenue does not always come with the optimal efficiency, as illustrated in Fig. 8d. Although enrolling more buyers increases the number of potential transaction pairs, it essentially lowers the admission threshold, making each transaction less profitable. However, high revenue is indeed extracted when the number of winners becomes a dominating factor.

7.2 District-D

Though *District-U* does not require *a priori* information and is *ex post* budget balanced, its efficiency is highly constrained due to its trade-reduction nature: many feasible trades are reduced to avoid a budget deficit. As shown in Fig. 8e, when the bid distribution knowledge is available, *District-D* can do better. As more sellers become available and the channel supplies increase, the auction efficiency can grow quickly, until the market is saturated with almost all buyers being served.

Interestingly, as shown in Fig. 8e, the more buyers join the market, the higher the mean auction efficiency is. We give an intuitive explanation here. The design rationale of *District-D* is to add as many trades as possible, with the constraint that the total revenue measured by virtual valuations is non-negative in every iteration. However, not all transactions are profitable. These transactions are accepted only when the auctioneer has already accumulated sufficient revenue in the past trades: it can use its accumulation to compensate for the deficit caused by these unprofitable transactions while still maintaining budget balance. For a smaller market with fewer buyers, the accumulated revenue is limited and is insufficient to compensate for the deficit. As a result, the trade that is not profitable has to be dropped, leaving a relatively low auction efficiency compared with a market with more buyers. Fig. 8g validates this point of view — low efficiency usually comes with low revenue. As a conclusion, *District-D* is scalable to large networks.

Also, we depict the standard deviation of the auction efficiency in Fig. 8f. We see that the scenario containing more participants, either buyers or sellers, tends to have a smaller standard deviation. Moreover, the deviation is

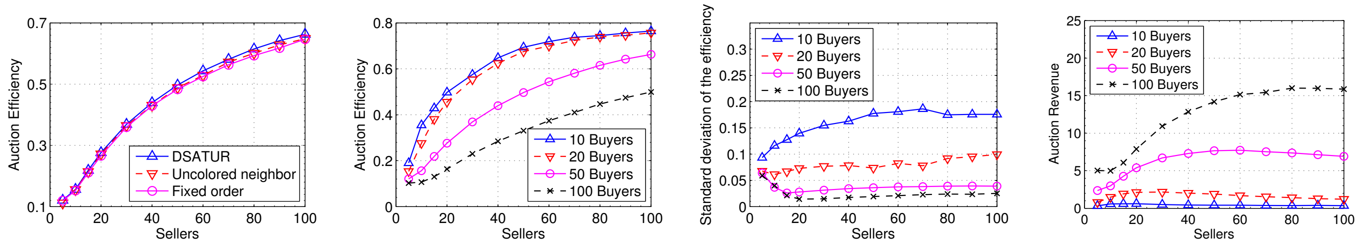

(a) *District-U* with N = 50: Mean auction efficiency with different *GraphColoring*().
(b) *District-U* with DSATUR: Mean auction efficiency.
(c) *District-U* with DSATUR: Standard deviation of the auction efficiency.
(d) *District-U* with DSATUR: Mean auction revenue.

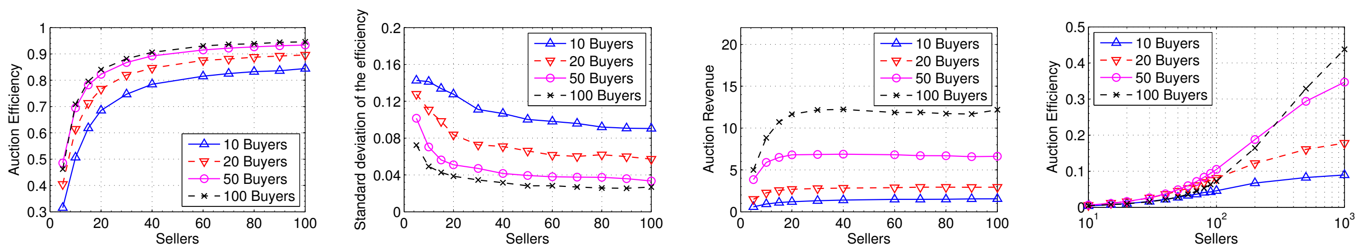

(e) *District-D*: Mean auction efficiency.
(f) *District-D*: Standard deviation of the auction efficiency.
(g) *District-D*: Mean auction revenue.
(h) TRUST-extension: Mean auction revenue.

Fig. 8. Simulation results of the efficiency-driven spectrum double auctions.

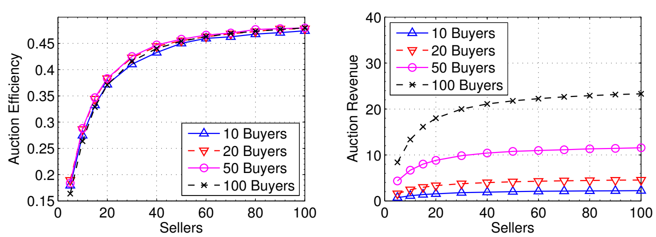

(a) Mean auction efficiency.
(b) Mean auction revenue.

Fig. 9. Simulation results for profit-driven *District-D*.

observed to be exponentially diminished with respect to the number of participants. One can view this result as an outcome of the Law of Large Numbers.

We next compare the performance of *District-D* and *District-U* with DSATUR by inspecting Fig. 8b, 8d, 8e and 8g. By a significant margin, *District-D* outperforms *District-U* in the auction efficiency, especially in a market containing more buyers.

Finally, we evaluate the profit-driven *Distric-D* discussed in Sec. 6.2. We see in Fig. 9b that the design generates significantly higher revenues than both *District-U* and its efficiency-driven counterpart. At a price, however, the auction efficiency is much lower, as illustrated in Fig. 9a.

In conclusion, *District-U* serves as an appropriate transitional mechanism for the auctioneer to sustain the auction without external subsidies. It generates a moderate level of auction efficiency, and creates time for the auctioneer to collect the distribution information of bids and asks. When the distribution information is available, the auctioneer can switch to *District-D* to enjoy higher auction efficiency.

7.3 District vs. A Simple Extension of TRUST

Sec. 6 discusses a scheme to directly extend an existing spectrum auction to make it compatible with local markets. A natural question is whether such a simple extension can provide acceptable performance. To study this, we extend TRUST[4] as described in Sec. 6 and investigate its auction efficiency.

In Fig. 8h, we see that, with market sizes comparable to those experimented in *District* (i.e., fewer than 100 sellers and buyers), the auction efficiency of TRUST-extension is fairly low — generally less than 0.1 — and grows slowly when channel supplies increase. Moreover, when more buyers are available in the auction, the market is saturated and the efficiency drops. This is in stark contrast to the scalability of *District-D*. The efficiency of TRUST-extension improves only when a very large amount of channel supplies are available in the market, but it is still severely limited when the number of buyers is small. By comparing Fig. 8h with both Fig. 8b and Fig. 8e, we conclude that *District* significantly outperforms the simple extension in Sec. 6.

8 Conclusion

In this paper, we present *District*, a set of new spectrum double auctions that incorporate market locality for practical spectrum markets, where sellers can freely partition their license areas to either sell or reserve, based on their own requirements. An auctioneer can start from *District-U*,

4. Similar to [5], the associated spectrum allocation algorithm is *Greedy-U*.

a uniform pricing auction, to obtain moderate auction efficiency without any *a priori* information about bids. After accumulating sufficient knowledge of bid distributions, it can then switch to *District-D*, a discriminatory pricing auction, to pursue better auction efficiency. Our computationally efficient designs are proved to be economically robust and scalable to large networks. With minor modifications, our designs can also be applied by a profit-driven auctioneer to pursue the high auction revenue. To our knowledge, this is the first set of spectrum double auctions designed for local markets with these properties.

REFERENCES

[1] W. Wang, B. Li, and B. Liang, "District: Embracing local markets in truthful spectrum double auctions," in *Proc. IEEE SECON*, 2011.
[2] FCC, "TV white spaces databases and database administrators," http://transition.fcc.gov/oet/whitespace/WhiteSpacesDAGuide.pdf.
[3] Spectrum Bridge, Inc., "SpecEx," http://www.specex.com.
[4] X. Zhou, S. Gandhi, S. Suri, and H. Zheng, "eBay in the sky: Strategy-proof wireless spectrum auctions," in *Proc. ACM MobiCom*, 2008.
[5] X. Zhou and H. Zheng, "TRUST: A general framework for truthful double spectrum auctions," in *Proc. IEEE INFOCOM*, 2009.
[6] J. Jia, Q. Zhang, Q. Zhang, and M. Liu, "Revenue generation for truthful spectrum auction in dynamic spectrum access," in *Proc. ACM MobiHoc*, 2009.
[7] H. Xu and B. Li, "A secondary market for spectrum," in *Proc. IEEE INFOCOM*, 2010.
[8] S. Wang, P. Xu, X. Xu, S. Tang, X. Li, and X. Liu, "TODA: Truthful online double auction for spectrum allocation in wireless networks," in *Proc. IEEE DySpan*, 2010.
[9] A. Gopinathan, Z. Li, and C. Wu, "Strategyproof auctions for balancing social welfare and fairness in secondary spectrum markets," in *Proc. IEEE INFOCOM*, 2011.
[10] R. P. McAfee, "A dominant strategy double auction," *J. Econ. Theory*, vol. 56, no. 2, pp. 434–450, 1992.
[11] R. B. Myerson, "Optimal auction design," *Math. Oper. Res.*, vol. 6, no. 1, p. 58, 1981.
[12] V. Krishna, *Auction Theory*. Academic Press, 2009.
[13] J. McMillan, "Selling spectrum rights," *J. Econ. Perspect.*, vol. 8, no. 3, pp. 145–162, 1994.
[14] R. McAfee and J. McMillan, "Analyzing the airwaves auction," *J. Econ. Perspect.*, vol. 10, no. 1, pp. 159–175, 1996.
[15] P. Cramton, "The FCC spectrum auctions: An early assessment," *J. Econ. Manage. Strategy*, vol. 6, no. 3, pp. 431–495, 1997.
[16] P. Cramton and J. Schwartz, "Collusive bidding: Lessons from the FCC spectrum auctions," *J. Regul. Econ.*, vol. 17, no. 3, pp. 229–252, 2000.
[17] J. Huang, R. Berry, and M. Honig, "Auction mechanisms for distributed spectrum sharing," in *Proc. Allerton Conf.*, 2004.
[18] M. Buddhikot and K. Ryan, "Spectrum management in coordinated dynamic spectrum access based cellular networks," in *Proc. IEEE DySpan*, 2005.
[19] S. Gandhi, C. Buragohain, L. Cao, H. Zheng, and S. Suri, "A general framework for wireless spectrum auctions," in *Proc. IEEE DySpan*, 2007.
[20] A. Subramanian, M. Al-Ayyoub, H. Gupta, S. Das, and M. Buddhikot, "Near-optimal dynamic spectrum allocation in cellular networks," in *Proc. IEEE DySpan*, 2008.
[21] F. Wu and N. Vaidya, "Small: A strategy-proof mechanism for radio spectrum allocation," in *Proc. IEEE INFOCOM*, 2011.
[22] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*. Cambridge, U.K.: Cambridge Univ. Press, 2007.
[23] R. B. Myerson and M. A. Satterthwaite, "Efficient mechanisms for bilateral trading," *J. Econ. Theory*, vol. 29, no. 2, pp. 265–281, 1983.
[24] A. Archer and É. Tardos, "Truthful mechanisms for one-parameter agents," in *Proc. IEEE FOCS*, 2001.
[25] D. Brélaz, "New methods to color the vertices of a graph," *Commun. ACM*, vol. 22, no. 4, pp. 251–256, 1979.

Wei Wang received the B.Engr. and M.A.Sc degrees from the Department of Electrical Engineering, Shanghai Jiao Tong University, China. Since 2010, he has been a Ph.D. candidate in the Department of Electrical and Computer Engineering at the University of Toronto. His research interests include cloud computing, network economics, and wireless networking. His current focus is on the design and analysis of market-driven resource allocations in cloud computing, using economics and optimization theory. He has also worked on spectrum sharing in cognitive radio networks using game theory and mechanism design.

Ben Liang received honors-simultaneous B.Sc. (valedictorian) and M.Sc. degrees in Electrical Engineering from Polytechnic University (NYU-Poly) in Brooklyn, New York, in 1997 and the Ph.D. degree in Electrical Engineering with Computer Science minor from Cornell University in Ithaca, New York, in 2001. In the 2001 – 2002 academic year, he was a visiting lecturer and post-doctoral research associate at Cornell University. He joined the Department of Electrical and Computer Engineering at the University of Toronto in 2002, where he is now a Professor. His current research interests are in mobile networking and multimedia systems. He received an Intel Foundation Graduate Fellowship in 2000 and an Ontario MRI Early Researcher Award (ERA) in 2007. He was a co-author of the Best Paper Award at the IFIP Networking conference in 2005. He is an editor for the IEEE Transactions on Wireless Communications and an associate editor for the Wiley Security and Communication Networks journal. He serves on the organizational and technical committees of a number of conferences each year. He is a senior member of IEEE and a member of ACM and Tau Beta Pi.

Baochun Li received the B.Engr. degree from the Department of Computer Science and Technology, Tsinghua University, China, in 1995 and the M.S. and Ph.D. degrees from the Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, in 1997 and 2000. Since 2000, he has been with the Department of Electrical and Computer Engineering at the University of Toronto, where he is currently a Professor. He holds the Nortel Networks Junior Chair in Network Architecture and Services from October 2003 to June 2005, and the Bell Canada Endowed Chair in Computer Engineering since August 2005. His research interests include large-scale multimedia systems, cloud computing, peer-to-peer networks, applications of network coding, and wireless networks. Dr. Li was the recipient of the IEEE Communications Society Leonard G. Abraham Award in the Field of Communications Systems in 2000. In 2009, he was a recipient of the Multimedia Communications Best Paper Award from the IEEE Communications Society, and a recipient of the University of Toronto McLean Award. He is a member of ACM and a senior member of IEEE.