

District: Embracing Local Markets in Truthful Spectrum Double Auctions

Wei Wang, Baochun Li and Ben Liang
Department of Electrical and Computer Engineering
University of Toronto

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 in *District* are tailored to cases with and without knowledge of bid distributions. An auctioneer can start from one design without any *a priori* information, and then switch to the other alternative after accumulating sufficient distribution knowledge. Complementary simulation studies show that spectrum utilization can be significantly improved when distribution information is available.

I. 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 channel access.

Yet, it is important to point out that transactions take place in secondary markets where spectrum is leased in a *local geographical region*. Unlike physical commodities that can be traded all over the world, spectrum is a local resource and is available only in *local markets* — only within the license region can a user be able to access the channel.

Practical spectrum markets, such as *SpecEx* [1], take advantage of this locality and provide flexible selling options to attract participation. They allow license holders to partition their license areas and decide which pieces to sell and which

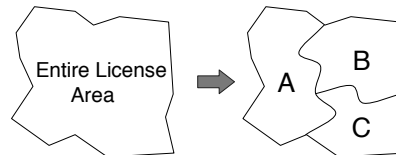


Fig. 1: A license holder partitions its entire license area into three regions, A, B and C. It can sell any of the pieces in the local spectrum market.

pieces to reserve. As suggested by the FCC, the holder uploads her decisions — the vacant spectrum as well as the associated geographical area — to an Internet database [2]. An unlicensed wireless device then queries the database to obtain a list of channels that can be used at the device’s location. Fig. 1 illustrates an example, where the license holder divides the entire license area into three regions and can sell any of the pieces to wireless users.

Unfortunately, market locality, as an inherent characteristic of spectrum markets in practice, is seldom mentioned in the literature. Most existing spectrum auctions [3], [4], [5], [6], [7], [8], [9], [10] are designed based on *global markets*, where channels are globally accessible to all users, and have to be traded as a whole in the entire license area. This impractical assumption seriously degrades 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 limited local regions, rather than the entire license area of a license holder. Wireless users who are outside of the limited local regions are not able to access these available channels. For example, in Fig. 1, a channel may be vacant in region A, yet utilized by its license holder in region B and C. In global markets, however, a license holder has to decide if it is able to make channels vacant in its entire license area, and partially available channels that are only vacant in some of the regions are not ready for sale.

To bridge such a gap between existing literature and practical limitations of geographical spectrum locality, in this paper, we present *District*, a set of new spectrum double auction mechanisms 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 used 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 [4], [6], [7] are based on McAfee’s design [11], which is for global markets only. Their direct extensions, unfortunately, is 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* [12] to double auctions and designs a market that price discriminates. In the latter case, *District* is designed to price uniformly. Both designs are proved to be computationally efficient and economically robust. Extensive simulation studies show that *District* substantially improves spectrum utilization with local markets, and is scalable to large networks.

The remainder of this paper is organized as follows. In Sec. II, we present the system model and formulate the problem considered in this paper. In Sec. III, we show that simple extensions to existing spectrum auction designs are not feasible. Sec. V and Sec. IV formally present the two designs of *District* in cases without and with distribution information, respectively. Extensive simulation results are given in Sec. VI. Sec. VII reviews related work and Sec. VIII concludes the paper.

II. SYSTEM MODEL

In a practical spectrum market (e.g., *SpecEx*), license holders sell the rights to access their under-used channels, while wireless users attempt to buy channel access at affordable prices. This can be modeled by a double auction with one auctioneer. In each round, every seller has one channel for sale in an indicated license area — called the *local market* — in which the channel is vacant, e.g., region A in Fig. 1. Every seller reports the channel, the associated local market, and an *ask* to the auctioneer, while every buyer requests to buy one channel by submitting a *bid* 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 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 constraint 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 local markets are drawn as circles only for illustration. In fact, they can have any shape and may not even be contiguous in general. We assume there are M participating license holders and N wireless users.

A. Modelling Channel Transactions within Local Markets

Channels should be assigned without introducing interference. We use a conflict graph $G = (V, E)$ to represent the

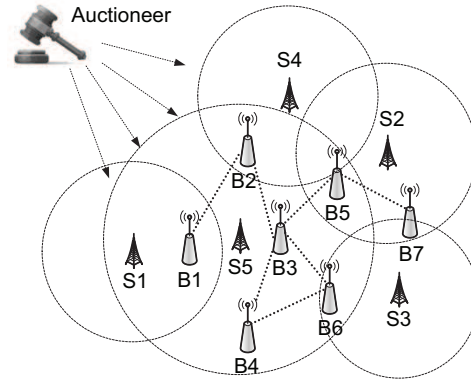


Fig. 2: A spectrum double auction with 7 buyers and 5 sellers. An auctioneer performs the auction among sellers and buyers. Sellers can partition their license areas and sell any pieces of their spectrum in the local market. All license areas for sale are circular in this figure, but can have any shape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

interference relation 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. In our example shown in Fig. 2, seven conflicting pairs of buyers are illustrated with 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 . The region of the local market is defined by m . The set of tradable sellers of buyer n is denoted as C_n . For example, in Fig. 2, B_1 ’s tradable sellers are S_1 and S_5 , i.e., $C_{B_1} = \{S_1, S_5\}$.

A channel assignment scheme is *feasible* if all transactions are 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 *graph coloring* scheme by treating C_n as the available colors that can be used to color a node n in G . In this sense, a buyer n is assigned to a tradable seller m if node $n \in V$ is colored by $m \in C_n$, and vice versa. As an example, one feasible channel assignment in Fig. 2 is $\{(B_1, S_1), (B_2, S_4), (B_4, S_5), (B_5, S_2), (B_6, S_3), (B_7, S_3)\}$. 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 $C = \{C_n | n \in V\}$.

B. 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 the bid submitted by 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. a_m^t (b_n^t) is the private information of seller m (buyer n), and is unknown to anyone else (including the auctioneer). Note that the seller m may submit a different value other than the true ask (i.e., $a_m \neq$

a_m^t), as long as it believes that this is more beneficial. The same may also be adopted by buyer n (i.e., $b_n \neq b_n^t$).

After collecting all asks $\mathbf{a} = (a_1, \dots, a_M)$ and bids $\mathbf{b} = (b_1, \dots, b_N)$, the auctioneer clears the market by computing the assignment of channels and winner payoffs. The assignment is represented by the coloring of conflict graph G , as mentioned above. Every winning seller m is paid p_m for leasing a channel, while every winning buyer n is charged c_n , both by the auctioneer. Therefore, the payoffs consist of both the payments $\mathbf{p} = (p_1, \dots, p_M)$ to sellers and the charges $\mathbf{c} = (c_1, \dots, c_N)$ 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 payment, charges, and corresponding utilities are zero. Also, the auctioneer gains a revenue, defined as the difference between the total charges and total payments, $\gamma = \sum_n c_n - \sum_m p_m$.

C. Economic Requirements

To encourage participation, an auction should satisfy some basic economic requirements [13] as defined below.

1) *Individual Rationality*: An auction is individually rational if no winning buyer is charged higher than its bid ($c_n \leq b_n$), and no winning seller is paid less than its ask ($p_m \geq a_m$). With this property, participants will always benefit by joining the auction.

2) *Budget Balance*: To make the auction self-sustained without any external subsidies, the generated revenue is required to be non-negative. Formally, an auction has *ex post* budget balance if $\gamma \geq 0$. A weaker requirement is *ex ante* budget balance, where the revenue is non-negative in expectation, i.e., $\mathbf{E}\gamma \geq 0$.

3) *Truthfulness*: Selfish participants can strategically bid to manipulate the market and obtain favorable outcomes by hurting the others. To avoid such manipulation, we should design a mechanism to ensure that all participants bid truthfully. Formally, an auction is truthful if no one can expect more benefit by cheating. That is, for all n (m), $b_n = b_n^t$ ($a_m = a_m^t$) is always the best bid (ask) with the maximum utility u_n^b (u_m^s), no matter how other participants behave.

We say an auction is *economically robust* [4], [7] if it is individually rational, budget balanced (either *ex post* or *ex ante*), and truthful.

D. Problem Definition

The motivation for introducing a spectrum market is to improve channel utilization. We therefore prefer facilitating as many wireless users as possible to access idle channels. For an auction mechanism \mathcal{M} , we define *auction efficiency* as the proportion of winning buyers:

$$\eta_{\mathcal{M}} = \frac{N_w}{N}, \quad (1)$$

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 . Ideally, we would like to find an economically robust

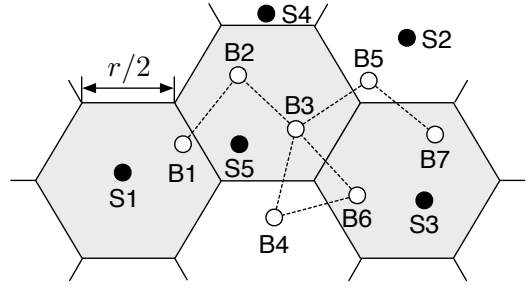


Fig. 3: A simple extension to TRUST [4].

auction mechanism that also maximizes the auction efficiency. However, the *impossibility theorem* [14] dictates that 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:

$$\begin{aligned} \max_{\mathcal{M}} \quad & \eta_{\mathcal{M}} ; \\ \text{s.t.} \quad & \mathcal{M} \text{ is economically robust.} \end{aligned} \quad (2)$$

The proposed auction mechanisms in *District* are shown to be economically robust while maintaining high auction efficiency.

III. DESIGNING AUCTIONS FOR LOCAL MARKETS

To design auction mechanisms for local spectrum markets, we begin by investigating into the possibility of extending existing spectrum auctions based on global markets. We will see in the following that simple extensions are either not feasible or leading to fairly inefficient outcomes.

In single-sided auctions, only one side (either buyers or sellers) has bidding strategies. VERITAS, proposed in [3], is known as the first single-sided spectrum auction designed to be truthful. However, it is shown that when extended to double auctions the truthfulness no longer holds [4]. Similar challenges also apply to [5], [9], and [10]. Therefore, such extensions do not satisfy our requirement for economic robustness.

When double auctions are considered, there exist some works in the literature on their designs in a global spectrum market [4], [6], [7]. In [6] and [7], spectrum reuse is not considered, and a complete conflict graph is assumed. Neither of them is applicable to our system model.

In [4], an economically robust double auction design has been proposed, referred to as TRUST, in a global spectrum market. It is possible to propose a simple extension to TRUST for local markets if the traded license areas are of some special shape. For example, suppose all license areas are circular. As shown in Fig. 3, we can partition the entire geographical 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. In other words, within one cell, the market is global. The proposed extension applies TRUST in every cell to cover the whole market.

However, such an extension is problematic. First, the traded areas need to be of some special shapes, which is usually not the case in practice. Second, buyers are only allowed to trade with sellers whose area centers are within the same cell. As a result, many originally feasible transactions are blocked. For example, in Fig. 3, buyer $B2$ cannot trade with seller $S4$ even if $B2$ is within $S4$'s license area. For this reason, the auction efficiency can be fairly low, especially when the cell edge is short. We later verify this by our numerical results, shown in Sec. VI.

The failure to simply extend auction mechanisms designed for global markets indicates that a new design tailored to local markets is required. In this paper, we present *District*, 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.

IV. *District-U*: AUCTION WITH UNIFORM PRICING

District-U adopts uniform pricing policies, such that all buyers (sellers) are charged (paid) exactly the same 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 explicitly reduced, 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.

A. Preliminaries

Without loss of generality, we assume bids and asks are sorted, from the most competitive to the least competitive:

$$b_1 \geq b_2 \geq \dots \geq b_N \geq 0; \quad (3)$$

$$0 \leq a_1 \leq a_2 \leq \dots \leq a_M. \quad (4)$$

For convenience, we introduce the following notations.

- $G_{M',N'}$: A subgraph of $G = (V, E, C)$ with only the first M' colors (sellers) and the first N' nodes (buyers). That is, $G_{M',N'} = (V', E', C')$ where $V' = \{1, \dots, N'\}$ and $C' = \{C'_1, \dots, C'_N\}$ where $C'_j = C_j \cap \{1, \dots, M'\}$.
- $\text{GraphColoring}(G)$: A graph coloring algorithm returning a colored G . *District-U* accepts only a *deterministic* graph coloring algorithm with no randomness introduced.

B. Mechanism Design

To achieve trade reduction, *District-U* first decides how many buyers to admit. To this end, we require the auctioneer to pick a predefined N' , announcing that only the top N' buyers are admitted to continue with the auction, while the rest are dropped. The auctioneer then computes the set of M' sellers admitted, and subsequently determines the winners and payoffs. The details are given in Algorithm 1.

Algorithm 1 *District-U* for a predefined N'

1. $M' = \arg \max_{0 \leq i < M} a_{i+1} \leq b_{N'+1}$
 2. $\text{GraphColoring}(G_{M',N'})$
 3. Seller m trades with buyer n if node n is colored by m
 4. **for** each winning buyer n **do**
 5. $c_n = b_{N'+1}$
 6. **end for**
 7. **for** each winning seller m **do**
 8. $p_m = a_{M'+1}$
 9. **end for**
 10. **return** transactions and payoffs (\mathbf{c} and \mathbf{p}).
-

C. Economic Properties

District-U is economically robust for any *deterministic* graph coloring algorithm. To see this, we first prove that it is *bid monotonic* with *critical payoffs*, which leads to truthfulness and individual rationality [13]. We then prove that it is also *ex post* budget balanced.

Proposition 1: *District-U* is bid monotonic. That is, 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.

Proof: Suppose buyer n wins by bidding b_n . Then $b_{N'+1} \leq b_n < b'_n$. Hence, by bidding b'_n , n remains among the top N' buyers. It will be admitted again, and $G_{M',N'}$ remains unchanged. With the same $G_{M',N'}$ as the input, the output of GraphColoring remains the same (because it is deterministic), which implies that n wins again. The proof for seller m is similar. ■

Definition 1: For winning buyer n (seller m), we say b_n (a_m) is *critical* if n (m) wins by submitting $b'_n > b_n$ ($a'_m < a_m$) and loses by submitting $b'_n < b_n$ ($a'_m > a_m$), given the others' submissions remain unchanged.

Proposition 2: In *District-U*, $b_{N'+1}$ ($a_{M'+1}$) is critical for winning buyer n (seller m).

Proof: For winning buyer n , it will remain in the set of top N' buyers as long as it bids $b_n > b_{N'+1}$, which ensures it to be admitted and win out. However, if n bids $b_n < b_{N'+1}$, it will be replaced by buyer $N' + 1$ and lose.

For winning seller m , it will remain in the set of top M' sellers as long as it asks $a_m < a_{M'+1}$, which ensures it to be admitted and win out. However, suppose m asks $a_m > a_{M'+1}$ and wins. Then there must be at least one seller m' satisfying $a_m \leq a_{m'} \leq b_{N'+1}$, which implies that $a_{M'+1} < a_{m'} \leq b_{N'+1}$. But this is impossible due to the definition of M' (line 1 of Algorithm 1). Therefore m loses. ■

As shown in [13], a bid monotonic auction is truthful and individually rational if it always charges critical bids from buyers and pays critical asks to sellers. Therefore, Propositions 1 and 2 are sufficient for *District-U* to be truthful and individually rational.

Proposition 3: *District-U* is *ex post* budget balanced.

Proof: For every transaction made between seller m and buyer n , $p_m = a_{M'+1} \leq b_{N'+1} = c_n$. Since the number

of winning buyers is greater than or equal to the number of winning sellers, we have $\gamma = \sum_n c_n - \sum_m p_m \geq 0$. ■

Hence, we conclude as the follows.

Theorem 1: District-U is economically robust.

D. Performance and Complexity

The performance of *District-U* is dominated by the predefined N' and the adopted graph coloring algorithm.

On the one hand, increasing N' admits more buyers to join, which potentially improves the auction efficiency. On the other hand, increasing N' results in a lower cut-off bid $b_{N'+1}$ and a lower cut-off ask $a_{M'+1}$. Therefore, fewer sellers are admitted, which hurts the auction efficiency. It is worth mentioning that, if the auctioneer were to enumerate N' in an attempt to adaptively optimize the auction efficiency, that would break bid monotonicity, making the auction untruthful. Therefore, a fixed N' must be chosen *a priori*. Without any statistics about the possible bids and asks, one reasonable choice may be the midpoint, i.e., $N' = \lfloor N/2 \rfloor$. We later present numerical data that shows how the auction efficiency varies with different choices of N' .

As for the graph coloring algorithm, it not only affects the auction efficiency, but also dominates the computational complexity. Fortunately, *District-U* imposes no restrictions on the algorithm except being deterministic. Therefore, any coloring algorithm with high computational efficiency could be directly adopted by *District-U*.

In conclusion, *District-U* is designed as a suitable starting mechanism for auctioneers without *a priori* information on bids or asks. Its auction efficiency is upper bounded by $\bar{\eta} = N'/N$ due to the trade reduction nature. We will see in the following that, after accumulating sufficient bid and ask information, the auctioneer can enjoy higher efficiency by switching to *District-D*.

V. District-D: AUCTION WITH DISCRIMINATORY PRICING

When bid and ask distributions are available, one can expect higher efficiency via *District-D*, an auction that price discriminates (i.e., winners have different payoffs). We begin by using Myerson's *virtual valuations* [12] to characterize the expected revenue of a truthful spectrum auction \mathcal{M} . We show that an economically robust \mathcal{M} is equivalent to a weighted graph coloring with a non-negative expected sum weight. We design *District-D* based on this and show that the design is computationally efficient.

A. Preliminaries

For every buyer n , denote its bid distribution function by $F_n^b(\cdot)$ and the corresponding density function by $f_n^b(\cdot)$. The functions $F_m^s(\cdot)$ and $f_m^s(\cdot)$ are similarly defined for 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 [12], 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 $\psi_m(a_m)$ and

$\phi_n(b_n)$ as the virtual valuations of seller m asking for a_m and buyer n bidding b_n , respectively, where

$$\begin{aligned}\psi_m(a_m) &= a_m + \frac{F_m^s(a_m)}{f_m^s(a_m)}, \\ \phi_n(b_n) &= b_n - \frac{1 - F_n^b(b_n)}{f_n^b(b_n)}.\end{aligned}$$

We assume *regular* distributions [12], i.e., all $\phi_n(\cdot)$ and $\psi_m(\cdot)$ are increasing functions. Let $\mathbf{v} = (a_1, \dots, a_M, b_1, \dots, b_N)$ be the vector of submitted asks and bids. Denote by $\gamma_{\mathcal{M}}(\mathbf{v}, G)$ the revenue of auction \mathcal{M} with submissions \mathbf{v} and 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.

Lemma 1 (Extended Myerson): Given submissions \mathbf{v} , let $x_n(\mathbf{v}) = 1$ ($y_m(\mathbf{v}) = 1$) if n (m) wins, i.e., n is colored (m is used for coloring), and $x_n(\mathbf{v}) = 0$ ($y_m(\mathbf{v}) = 0$) otherwise. Then we have

$$\mathbf{E}_{\mathbf{v}}[\gamma(\mathbf{v})] = \mathbf{E}_{\mathbf{v}}\left[\sum_n \phi_n(b_n)x_n(\mathbf{v}) - \sum_m \psi_m(a_m)y_m(\mathbf{v})\right]. \quad (5)$$

The proof of the above lemma is similar to Myerson's theorem [12] and is omitted here. Lemma 1 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 bidding 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 . Then the right hand side of (5) can be defined as the *expected sum weight* of a colored graph G . Therefore, achieving *ex ante* budget balance is equivalent to maintaining a non-negative expected sum weight. Furthermore, we see that the expected revenue of a truthful mechanism can be totally characterized by winners only, independent of their payoffs. In other words, it suffices to focus only on winner-determination designs (i.e., graph coloring) to achieve *ex ante* budget balance.

B. Winner Determination

The first component of *District-D* is an algorithm for the auctioneer to determine the winning buyers and sellers.

In the proposed heuristic algorithm, at every iteration, we pick a feasible buyer-seller pair with the maximum marginal revenue measured by virtual valuation. If the total revenue is non-negative after adding the pair's marginal revenue, the pair is accepted. Otherwise, the pair is rejected and the algorithm terminates. For convenience, we introduce the following notations before presenting the formal algorithm in Algorithm 2.

- \mathcal{T} – Round-by-round record of transactions already made by the winner determination algorithm.
- \mathcal{T}^b – Set of winning buyers associated with \mathcal{T} .
- \mathcal{T}^s – Set of winning sellers associated with \mathcal{T} .

- $\Delta_{m,n}(\mathcal{T}, a_m, b_n)$ – Marginal revenue generated by assigning m to n , given \mathcal{T} , m 's ask a_m , and n 's bid b_n . By (5), we have

$$\Delta_{m,n}(\mathcal{T}, a_m, b_n) = \phi_n(b_n) \cdot I_{n \notin \mathcal{T}^b} - \psi_m(a_m) \cdot I_{m \notin \mathcal{T}^s}, \quad (6)$$

where $I_x = 1$ if condition x is satisfied, and $I_x = 0$ otherwise.

- *MaxMarginalRev*(\mathcal{T}) – Given \mathcal{T} , calculate the transaction (m, n) with the maximum marginal revenue among all feasible transactions, i.e., $\Delta_{m,n}(\mathcal{T}, a_m, b_n) = \max_{i,j} \Delta_{i,j}(\mathcal{T}, a_i, b_j)$. The return is (m, n, Δ) , where $\Delta = -\infty$ if no transaction is feasible, and $\Delta = \Delta_{m,n}(\mathcal{T}, a_m, b_n)$ otherwise.

Algorithm 2 *District-D* Winner Determination

1. **Initialization:** $\gamma \leftarrow 0$, $\mathcal{T} \leftarrow \emptyset$, and *stop* \leftarrow *false*.
 2. **while** *stop* = *false* **do**
 3. $(m, n, \Delta) \leftarrow$ *MaxMarginalRev*(\mathcal{T})
 4. **if** $\gamma + \Delta \geq 0$ **then**
 5. $\gamma \leftarrow \gamma + \Delta$
 6. **Make a deal:** Add (m, n) to \mathcal{T} , color n by m
 7. **else**
 8. *stop* \leftarrow *true*
 9. **end if**
 10. **end while**
 11. **return** \mathcal{T}
-

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 always colors G with a non-negative sum weight. Hence $\mathbb{E}[\gamma_{\mathcal{M}}(\mathbf{v}, G)] \geq 0$. Since this statement holds for every G , we conclude the proof. ■

Proposition 5: Algorithm 2 is bid monotonic.

Proof: We prove the buyer's case by contradiction. Suppose by submitting b_n (Bid 1), n wins in the k th iteration of Algorithm 2, while by submitting $b'_n > b_n$ (Bid 2), n loses.

For convenience, for Bid 2, denote the other buyers' bids by $b'_j = b_j, j \neq n$. Also denote $\Delta\phi_n = \phi_n(b'_n) - \phi_n(b_n)$. We see that $\Delta\phi_n \geq 0$ due to the assumption of an increasing $\phi_n(\cdot)$. By (6), we have:

$$\Delta_{i,j}(\mathcal{T}, a_i, b'_j) - \Delta_{i,j}(\mathcal{T}, a_i, b_j) = \Delta\phi_n \cdot I_{j=n}. \quad (7)$$

Let $(\mathcal{T}_l, \gamma_l)$ and $(\mathcal{T}'_l, \gamma'_l)$ be the vectors containing the transactions and total revenue after the l th iteration of Algorithm 2 with Bid 1 and Bid 2, respectively. Since buyer n does not win in the first $k-1$ iterations in either case, $\mathcal{T}_l = \mathcal{T}'_l, \gamma_l = \gamma'_l, l = 0, \dots, k-1$ ($\mathcal{T}_0 = \mathcal{T}'_0 = \emptyset, \gamma_0 = \gamma'_0 = 0$). Now for any feasible transaction pair (i, j) in the k th iteration with Bid 2, its marginal revenue is

$$\begin{aligned} \Delta_{i,j}(\mathcal{T}'_{k-1}, a_i, b'_j) &= \Delta_{i,j}(\mathcal{T}_{k-1}, a_i, b'_j) \\ &= \Delta_{i,j}(\mathcal{T}_{k-1}, a_i, b_j) + \Delta\phi_n I_{j=n}, \end{aligned} \quad (8)$$

where the second equality holds because of (7).

For Bid 1, suppose n trades with m in the k th iteration.

Then, (m, n) is of the maximum marginal revenue and maintains the budget balance, i.e.,

$$\Delta_{m,n}(\mathcal{T}_{k-1}, a_m, b_n) = \max_{i,j} \Delta_{i,j}(\mathcal{T}_{k-1}, a_i, b_j), \quad (9)$$

$$\gamma_{k-1} + \Delta_{m,n}(\mathcal{T}_{k-1}, a_m, b_n) \geq 0. \quad (10)$$

Now inspect the marginal revenue of the same transaction pair (m, n) in the k th round with Bid 2. Let $i = m$ and $j = n$. By (8), we have:

$$\begin{aligned} \Delta_{m,n}(\mathcal{T}'_{k-1}, a_m, b'_n) &= \Delta_{m,n}(\mathcal{T}_{k-1}, a_m, b_n) + \Delta\phi_n \\ &= \max_{i,j} \Delta_{i,j}(\mathcal{T}_{k-1}, a_i, b_j) + \Delta\phi_n \\ &= \max_{i,j} \Delta_{i,j}(\mathcal{T}'_{k-1}, a_i, b'_j). \end{aligned}$$

Here, the second equality holds because of (9); while the third equality follows from the fact that $\Delta\phi_n \geq 0$.

Thus, we conclude that (m, n) should also be selected in the k th iteration in Bid 2, because it generates the maximum marginal revenue and keeps the budget balanced. This contradicts the assumption that n loses with Bid 2.

With a similar argument, we see that the statement also holds for the seller's case. ■

Next, we present the pricing mechanisms, which find the critical bid and ask in a computationally efficient manner.

C. Buyer Pricing

By Definition 1, we see that the critical bid is the minimum submission required to win. The basic idea 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.

For simplicity, let \mathcal{T} be the transactions list generated by Algorithm 2. Let \mathcal{T}_l be the first l transactions in \mathcal{T} , i.e., $\mathcal{T}_l = \{(i_1, j_1), \dots, (i_l, j_l)\}$, where (i_l, j_l) is the l th transaction made between seller i_l and buyer j_l . Finally, let \mathcal{T}_l^s be 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., $c_n = b_n^l, l \leq k$. From Line 6, we see that $\phi_n(c_n) = \phi_n(b_n^l) = \phi_{j_l}(b_{j_l})$. For n bidding $b_n > c_n$, the worst case is that it loses in the first $l-1$ rounds ($l = 1, \dots, k$). But in the l th round, $\phi_n(b_n) > \phi_n(c_n) = \phi_{j_l}(b_{j_l})$, where the inequality holds since $\phi_n(x)$ is increasing. This implies that selecting n would generate more marginal revenue than selecting j_l , i.e., $\Delta_{i_l,n}(\mathcal{T}_{l-1}, a_{i_l}, b_n) > \Delta_{i_l,j_l}(\mathcal{T}_{l-1}, a_{i_l}, b_{j_l})$. Notice that (i_l, j_l) is of the maximum marginal revenue when n is absent, we conclude that (i_l, n) maximizes the marginal revenue. Therefore, n wins by being selected to trade with i_l .

Algorithm 3 *District-D* Buyer Pricing for winning buyer n

1. Remove n and run Algorithm 2 to obtain the transaction (*seller, buyer*) list $\mathcal{T} = \{(i_1, j_1), \dots, (i_k, j_k)\}$.
 2. $S_n \leftarrow n$'s tradable sellers, $c_n \leftarrow \infty$, and $\gamma \leftarrow 0$
 3. **for** $l = 1$ to k **do**
 4. $\gamma \leftarrow \gamma + \Delta_{i_l, j_l}(\mathcal{T}_{l-1}, a_{i_l}, b_{j_l})$
 5. **if** $i_l \in S_n$ **and** j_l conflicts with n **then**
 6. $b_n^l \leftarrow \phi_n^{-1}(\phi_{j_l}(b_{j_l}))$
 7. $c_n \leftarrow \min\{c_n, b_n^l\}$
 8. $S_n \leftarrow S_n \setminus \{i_l\}$
 9. **end if**
 10. **end for**
 11. **if** $S_n \neq \emptyset$ **then**
 12. $b_n^{k+1} \leftarrow \min_{i \in S_n} \phi_n^{-1}(\psi_i(a_i)I_{i \notin \mathcal{T}_k^s} - \gamma)$
 13. $c_n \leftarrow \min\{c_n, b_n\}$
 14. **end if**
 15. **return** c_n
-

Case 2: $c_n = b_n^{k+1}$. For n bidding $b_n > c_n$, the worst case is that it loses in the first k rounds. In this case, $\phi_n(b_n) > \phi_n(c_n) = \phi_n(b_n^{k+1}) = \min_{i \in S_n} \psi_i(a_i)I_{i \notin \mathcal{T}_k^s} - \gamma$. We assume $\psi_m(a_m)I_{m \notin \mathcal{T}_k^s} = \min_{i \in S_n} \psi_i(a_i)I_{i \notin \mathcal{T}_k^s}$ where $m \in S_n$. In this case, seller m can still trade with n after the first k rounds, with marginal revenue $\Delta_{m,n}(\mathcal{T}_k, a_m, b_n) = \phi_n(b_n) - \psi_m(a_m)I_{m \notin \mathcal{T}_k^s} > -\gamma$. By doing so, the total revenue remains non-negative, i.e., $\Delta_{m,n}(\mathcal{T}_k, a_m, b_n) + \gamma > 0$. Therefore, n wins.

Next, if n bids less than c_n (i.e., $b_n < c_n$), then it loses in the first k rounds: $\phi_n(b_n) < \phi_n(b_n^l) = \phi_{j_l}(b_{j_l})$ for all $l = 1, \dots, k$ implies smaller marginal revenue, i.e., $\Delta_{i_l, n}(\mathcal{T}_{l-1}, a_{i_l}, b_n) < \Delta_{i_l, j_l}(\mathcal{T}_{l-1}, a_{i_l}, b_{j_l})$. Moreover, even if $S_n \neq \emptyset$ after the first k rounds, n cannot trade with any seller $m \in S_n$, for the total revenue would be negative otherwise: $\Delta_{m,n}(\mathcal{T}_k, a_m, b_n) + \gamma = \phi_n(b_n) - (\psi_m(a_m)I_{m \notin \mathcal{T}_k^s} - \gamma) < 0$. Therefore, n loses. ■

D. Seller Pricing

For sellers, the analysis on pricing is similar. Seller m only needs to ask lower than its competitors to win the auction. We first remove m and run winner determination to see its winning competitors' asks. Seller m can win by asking lower than these competitors. The detailed procedure is given in Algorithm 4.

In Algorithm 4, we see that if m asks lower than a_m^l calculated in Line 6 ($a_m < a_m^l, l = 1, \dots, k$), it wins. The worst case for m is that it loses in the first $l-1$ rounds, but in the l th round, $\psi_m(a_m) < \psi_m(a_m^l) = \phi_{j_l}(b_{j_l}) - \Delta_{i_l, j_l}(\mathcal{T}_{l-1}, a_{i_l}, b_{j_l})$, which implies that assigning (m, j_l) would generate more marginal revenue than assigning (i_l, j_l) :

$$\begin{aligned} \Delta_{m, j_l}(\mathcal{T}_{l-1}, a_m, b_{j_l}) &= \phi_{j_l}(b_{j_l}) - \psi_m(a_m) \\ &> \Delta_{i_l, j_l}(\mathcal{T}_{l-1}, a_{i_l}, b_{j_l}). \end{aligned} \quad (11)$$

Therefore m is selected instead of i_l .

It is possible that even with m absent, there remain tradable buyers left unassigned. In other words, m does not have competitors for these buyers. To trade with any of these buyers,

Algorithm 4 *District-D* Seller Pricing for winning seller m

1. Remove m and run Algorithm 2 to obtain the transaction (*seller, buyer*) list $\mathcal{T} = \{(i_1, j_1), \dots, (i_k, j_k)\}$.
 2. $B_m \leftarrow m$'s tradable buyers, $p_m \leftarrow -\infty$, and $\gamma \leftarrow 0$
 3. **for** $l = 1$ to k **do**
 4. $\gamma \leftarrow \gamma + \Delta_{i_l, j_l}(\mathcal{T}_{l-1}, a_{i_l}, b_{j_l})$
 5. **if** j_l is m 's tradable buyer **then**
 6. $a_m^l \leftarrow \psi_m^{-1}(\phi_{j_l}(b_{j_l}) - \Delta_{i_l, j_l}(\mathcal{T}_{l-1}, a_{i_l}, b_{j_l}))$
 7. $p_m \leftarrow \max\{p_m, a_m^l\}$
 8. $B_m \leftarrow B_m \setminus \{j_l\}$
 9. **end if**
 10. **end for**
 11. **if** $B_m \neq \emptyset$ **then**
 12. $a_m^{k+1} \leftarrow \max_{j \in B_m} \psi_m^{-1}(\phi_j(b_j) + \gamma)$
 13. $p_m \leftarrow \max\{p_m, a_m\}$
 14. **end if**
 15. **return** p_m
-

m only needs to keep the total revenue non-negative, which is true if it asks for no more than a_m^{k+1} calculated in Line 12.

On the other hand, if m asks for more than any a_m^l , it loses — it either cannot beat some competitors or cannot maintain the budget balance. Therefore, we have:

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

E. Economic Properties and Computational Efficiency

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

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

Now we discuss 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. Within the loop, the complexity is dominated by $MaxMarginalRev(\cdot)$. A simple implementation is to go through all tradable transactions without conflicting with previously made trades, and to select the one with the maximum marginal revenue. There are at most MN such transactions, each requiring at most N comparisons to clarify the conflicting relation. We hence need $O(MN^2)$ time for $MaxMarginalRev(\cdot)$ and $O(MN^3)$ time for Algorithm 2. Note that Algorithm 2 also dominates the complexity of Algorithm 3 and 4, and runs at most N times to calculate the payoffs in each case. By omitting the constant factor, we conclude that *District-D* runs within $O(MN^4)$ time.

VI. 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. For simplicity, the license area is set to be circular, with radius uniformly distributed in $[0.2, 0.5]$ and

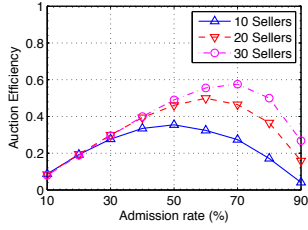


Fig. 4: Trade-off on auction efficiency in *District-U* with 50 buyers.

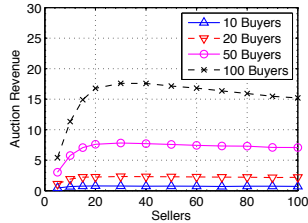


Fig. 6: Auction revenue of *District-U* with 50% buyers enrolled.

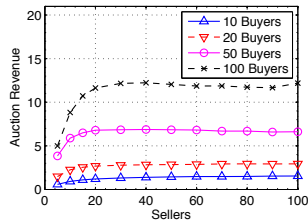


Fig. 8: Auction revenue of *District-D*.

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 true asks. We use a greedy graph coloring algorithm in *District-U* throughout the experiments. That is, we always color the node with the fewest uncolored neighbors.

Since *District* is already proven to be economically robust, we further evaluate its performance in terms of auction efficiency only. We also present numerical results for the generated revenue to study the relation between auction efficiency and revenue. However, since maintaining budget balance is sufficient for *District*, revenue maximization is outside the scope of this work, some in-depth discussions of which can be found in [5], [10]. Each result obtained below has been averaged over 10000 runs.

District-U. As a trade reduction auction, *District-U* only admits top buyers with high bids to join. Specifically, the admission rate is determined as $r = N'/N$, where N' is a predefined cut-off ranking for buyers. There is a trade-off in choosing r (or equivalently, N'), as discussed in Sec. IV. Fig. 4 illustrates such a trade-off. For a fixed number of buyers and sellers, neither high nor low admission rate is beneficial to auction efficiency, since in both cases the number of tradable

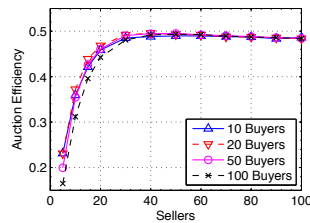


Fig. 5: Auction efficiency of *District-U* with 50% buyers enrolled.

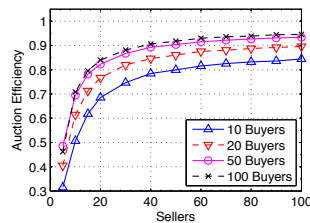


Fig. 7: Auction efficiency of *District-D*.

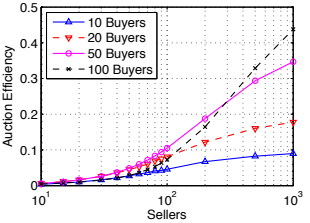


Fig. 9: Auction efficiency of TRUST-extension.

pairs is reduced. This shortage is gradually alleviated when more sellers (from 10 to 30 in the experiment) join the market to provide more channels. Also notice that the auction efficiency is a linear function of the admission rate when the rate is low, which implies that the low efficiency is caused by the deficit in channel supplies.

Next, we focus on the middle ground with $r = 50\%$, not only because it is a reasonable choice when no bid information is known *a priori*, but also because *District-U* only serves as a temporary solution for the auctioneer until sufficient bid knowledge is available. Fig. 5 and 6 illustrate the auction efficiency and revenue when $r = 50\%$. We see that the auction efficiency differs little in all four experiments, except a slight decrease with the growth of buyers. Also notice that the deficit of channel supplies dominates the auction efficiency until the number of sellers exceeds 20, after which the efficiency is constrained by the adopted admission rate. The market is saturated with stable efficiency with around 30 sellers, and the efficiency is upper bounded by 0.5. In addition, we see that the maximum revenue does not come with the optimal efficiency, but high revenue is indeed generated when the number of winners becomes a dominating factor.

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. 7, 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 without a hard cap constrained, until the market is saturated with almost all buyers served.

Interestingly, as shown in Fig. 7, the more buyers join the market, the higher the 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 stage. However, not all transactions are profitable. If the auction has accumulated sufficient revenue in the past trades, then it can compensate for the deficit caused by these transactions while still maintaining budget balance. For a smaller market with fewer buyers, the revenue accumulated 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. 8 validates this point of view — low efficiency usually comes with low revenue. In this sense, *District-D* is scalable to large networks.

We next compare the performance of *District-D* and *District-U* ($r = 50\%$) by inspecting Fig. 5, 6, 7 and 8. As expected, *District-D* outperforms *District-U* in auction efficiency, with significant efficiency gain in a market containing more buyers. It is interesting to see that for a smaller market with fewer buyers, *District-D* generates more revenue than *District-U*, although the gap is small. However, as more buyers join the market, the revenue generated by *District-U* overwhelms *District-D*. The reason behind this is that *District-U* always

makes profitable transactions, while *District-D* accepts unprofitable trades that are compensated with the accumulated revenue. As a result, *District-U* serves as an appropriate starting mechanism for the auctioneer to sustain the auction without external subsidies. With moderate auction efficiencies, it makes time for the auctioneer to collect bid distribution information. Once the distribution is available, the auctioneer can switch to *District-D* to pursue higher auction efficiency.

District vs. a Simple Extension of TRUST. Sec. III discusses a scheme to directly extend existing spectrum auctions into local markets. A natural question is whether such a simple extension provides acceptable performance. To study this, we extend TRUST (with *Greedy-U* as an allocation algorithm [4]) as described in Sec. III and investigate its auction efficiency. In Fig. 9, we see that, with market sizes comparable to those experimented in *District* (i.e., fewer than 100 sellers and buyers), the efficiency of TRUST-extension is fairly low (generally less than 0.1) and grows slowly when channel supplies increase. Moreover, when more buyers join 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. 9 with both Fig. 5 and Fig. 7, we conclude that *District* significantly outperforms the simple extension in Sec. III.

VII. RELATED WORK

Auction mechanisms serve as an efficient way to distribute scarce resources in a market. To encourage participation, the auction is required to be economically robust. Many well-known mechanisms, including both single-sided and double-sided auctions, are designed to achieve truthfulness for physical commodities [11], [12], [14], [15].

As a counterpart of the commodity auction, spectrum auction provides efficient solution to distribute spectrum resources in wireless networks. Early works include transmit power auctions [16] and spectrum band auctions [17], [18]. Truthfulness is first considered for single-sided spectrum auction in [3], where spectrum reuse is supported. Similar model is also adopted by the following work: [5] and [10] aim to generate maximum revenue for the seller; [9] supports channel reserve prices in the model; and [8] studies the fairness issue in spectrum allocations. Truthful spectrum double auction is first designed in [4]. [6] further advocates conducting double auctions in spectrum secondary markets, with participants being secondary users. [7] takes the time domain into account and proposes a truthful online auction. All these works discuss truthfulness in the sense of global markets, where the spectrum to be auctioned is globally accessible to all buyers and to be sold as a whole in the entire license area of a primary license holder.

However, market locality is essential in spectrum trading and is emphasized by the recent push for white spaces databases [2] maintained by a third party such as *Spectrum*

Bridge [1]. Spectrum owners should be allowed to partition their entire license area and lease parts of it to enterprises and consumers. *District* caters to their requirements and differs significantly from previous works by considering such market locality, while achieving economic robustness and efficient spectrum reuse.

VIII. 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*, 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 high auction efficiency. Our computationally efficient designs are proved to be economically robust and scalable to large networks. To our knowledge, this is the first set of double spectrum auctions designed for local markets with these properties.

REFERENCES

- [1] Spectrum Bridge, Inc., <http://www.specex.com>.
- [2] FCC, "TV white spaces databases and database administrators," <http://www.fcc.gov/oet/whitespace/tvwpdda.html>.
- [3] X. Zhou, S. Gandhi, S. Suri, and H. Zheng, "eBay in the sky: Strategy-proof wireless spectrum auctions," in *Proc. ACM MobiCom*, 2008, pp. 2–13.
- [4] X. Zhou and H. Zheng, "TRUST: A general framework for truthful double spectrum auctions," in *Proc. IEEE INFOCOM*, 2009, pp. 999–1007.
- [5] J. Jia, Q. Zhang, Q. Zhang, and M. Liu, "Revenue generation for truthful spectrum auction in dynamic spectrum access," in *Proc. ACM MobiHoc*, 2009, pp. 3–12.
- [6] H. Xu and B. Li, "A secondary market for spectrum," in *Proc. IEEE INFOCOM*, 2010.
- [7] 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.
- [8] A. Gopinathan, Z. Li, and C. Wu, "Strategyproof auctions for balancing social welfare and fairness in secondary spectrum markets," in *Proc. IEEE INFOCOM*, 2011, pp. 3020–3028.
- [9] F. Wu and N. Vaidya, "SMALL: A strategy-proof mechanism for radio spectrum allocation," in *Proc. IEEE INFOCOM*, 2011, pp. 81–85.
- [10] M. Al-Ayyoub and H. Gupta, "Truthful spectrum auctions with approximate revenue," in *Proc. IEEE INFOCOM*, 2011, pp. 2813–2821.
- [11] R. P. McAfee, "A dominant strategy double auction," *Journal of Economic Theory*, vol. 56, no. 2, pp. 434–450, 1992.
- [12] R. B. Myerson, "Optimal auction design," *Mathematics of Operations Research*, vol. 6, no. 1, pp. 58–73, 1981.
- [13] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [14] R. B. Myerson and M. A. Satterthwaite, "Efficient mechanisms for bilateral trading," *Journal of Economic Theory*, vol. 29, no. 2, pp. 265–281, 1983.
- [15] W. Vickrey, "Counterspeculation, auctions, and competitive sealed tenders," *The Journal of Finance*, vol. 16, no. 1, pp. 8–37, 1961.
- [16] J. Huang, R. Berry, and M. Honig, "Auction-based spectrum sharing," *Mobile Networks and Applications*, vol. 11, no. 3, pp. 405–418, 2006.
- [17] M. Buddhikot and K. Ryan, "Spectrum management in coordinated dynamic spectrum access based cellular networks," in *Proc. IEEE DySPAN*, 2005, pp. 299–307.
- [18] S. Gandhi, C. Buragohain, L. Cao, H. Zheng, and S. Suri, "A general framework for wireless spectrum auctions," in *Proc. IEEE DySPAN*, 2007, pp. 22–33.