Fair Pricing in the Sky: Truthful Frequency Allocation with Dynamic Spectrum Supply

Qihang Sun, Qian Wang*
School of Computer
Wuhan University
Wuhan 430072, Hubei, China
Email: {qianwang, qhsun}@whu.edu.cn

Kui Ren
Department of CSE
The State University of New York at Buffalo
Buffalo, NY 14260-2500, USA
Email: kuiren@buffalo.edu

Xiaohua Jia
Department of CS
City University of Hong Kong
Kowloon, Hong Kong
Email: csjia@cityu.edu.hk

***Abstract*—Spectrum auctions, which allow a spectrum owner to sell licenses for signal transmission over specific bands, can allocate scarce spectrum resources quickly to the users that value them most and have received a great deal of research attention in recent years. While enabling reusability-driven spectrum allocation, truthful spectrum auction designs are also expected to provide price fairness for homogeneous channels, online auction with unknown and dynamic spectrum supplies, and bounded system performance. Existing works, however, lack of such designs due to the inherent technically challenging nature. In this paper, we study the problem of allocating channels to spectrum users with homogeneous/heterogenous demands in a setting where idle channels are arriving dynamically, with the goal of maximizing social welfare. Taking spectrum reusability into consideration, we present a suite of novel and efficient spectrum auction algorithms that achieve fair pricing for homogeneous channels, strategy-proofness, online spectrum auction with a dynamic supply and a** log **approximation to the optimal social welfare. To the best of our knowledge, we are the first to design truthful spectrum auctions enabling fair payments for homogenous channels and spectrum reusability with dynamic spectrum supply. Experimental results show that our schemes outperform the existing benchmarks by providing almost perfect fairness of pricing for both single- and multi-unit demand spectrum users.**

I. Introduction

The rapid growth of wireless technologies and applications has increasingly made radio spectrum a critical yet scarce resource for wireless services. Traditional centralized and static spectrum allocations led to an inefficient use of spectrum resources, which motivates (*i.e.*, financial incentives) the design of market-based approaches for redistributing the idle spectrum, providing spectrum opportunities for unexploited licensed bands and gaining efficient spectrum utilization.

To allow a spectrum owner to sell licenses for signal transmission over specific bands, auctions can be used to allocate scarce spectrum resources quickly and efficiently to the users that value them most. In contrast to other digital goods, spectrum has a very unique characteristic called *reusability* due to the inherent nature of interference in radio transmissions. That is, users whose radio transmissions do not interfere each other in different geographic locations are able to share the same spectrum simultaneously. Obviously, spectrum/frequency reusability enables the communication system to increase both coverage and capacity. However, it also poses new challenges for the spectrum auction design, *e.g.*, the reusability makes it hard to achieve a truthful spectrum auction design.

To accommodate reusability-driven spectrum allocation, recently spectrum auctions have received extensive research efforts in the literature [1]–[7]. The key idea of addressing the problem of spectrum reusability is to divide bidders into multiple non-overlapping segments based on the interference constraints using graph coloring algorithms [8]. In those approaches, their design goals are to achieve spectrum auctions with either truthfulness [1], [3]–[7], or revenue maximization [6], or collusion-resistance [5], or privacy-preservation [4], or satisfaction of spectrum users with heterogeneous demands [7]. While theoretically sound, they only focus on the offline or say static auction model, where the set of users (*i.e.*, bidders) and the set of goods (*i.e.*, channels) are pre-determined before the start of auction process.

Recently, researches on online spectrum auction models have aroused much interest [9]–[11]. In [9], Deek et al. made an extension of [1] and investigated the online multi-good selling scenario. However, it does not provide a performance bound on revenue with respect to the optimal solution in general. Under the same online auction model, Xu et al. [10] proposed TOFU, another online semi-truthful spectrum auction scheme with channel *preemption*. Different from the previous solutions, TOFU achieves only semi-truthfulness, where users may be able to underbid to gain self-benefits. As a following work, Xu et al. [11] extended the results to multi-channel wireless networks. Almost without exception the above online designs consider the dynamic behaviors of spectrum users, where spectrum resources are fixed and spectrum users are arriving dynamically. In practice, however, the availability of spectrum resources are changing dynamically. That is, previously-occupied channels will be continually released and made available for unsatisfied spectrum users during one auction period. As far as we know, the existing spectrum auction solutions cannot be applied to this application scenario. Such online auction model with dynamically-arriving and unknown spectrum supplies has received limited research attention so far. In addition, almost all existing spectrum auction models consider the allocation of homogeneous channels with uniform characteristics, but the payments of different users for a channel vary greatly. From the perspective of users, it is unfair for them to make payments that differ considerably for the same goods, *i.e.*, homogeneous channels.

*Corresponding author.

978-1-4799-6204-4/14 $31.00 © 2014 IEEE
DOI 10.1109/ICNP.2014.45

To address the above concerns, in this paper we propose a suite of novel and efficient spectrum auction mechanisms. In particular, we introduce a two-level randomization into the selection of winning candidates and seamlessly integrate the channel allocation with carefully-designed pricing methods, achieving fair pricing for homogeneous channels, strategy-proofness, online auction with a dynamic spectrum supply and bounded performance. To the best of our knowledge, we are the first to design truthful spectrum auctions enabling fair pricing and spectrum reusability with dynamic spectrum supply. More specifically, our contributions are summarized as follows.

- We formulate and investigate the problem of allocating channels to spectrum users in a setting where idle channels are arriving dynamically and the total channel supply is unknown, with the goal of maximizing the social welfare.
- We propose a new and novel spectrum auction scheme for the single-unit demand case and show that it achieves all desirable properties, including truthfulness, price fairness, efficiency and online auctions with dynamic and unknown supply of idle channels.
- We extend our online spectrum auction system to support spectrum users with multi-unit demands. We prove that the proposed scheme again achieves truthfulness and show that *good* price fairness for each channel can still be achieved.
- We analytically show that the proposed schemes can achieve a $\log$ approximation to the optimal social welfare. Experimental results show that our schemes outperform the existing benchmarks by providing almost perfect price fairness for both single- and multi-unit demand spectrum users.

The rest of the paper is organized as follows. We present the preliminaries and definitions in Section II. We identify the design challenges, develop our spectrum auction mechanisms and provide theoretical analysis of their properties in Section III. In Section IV, we conduct experiments to evaluate and compare the performance of our spectrum auction mechanisms with the existing benchmarks. We discuss the related work in Section V and finally conclude our work in Section VI.

II. Problem Formulation and Preliminaries

Due to the scarcity of spectrum resource, we consider an online spectrum auction system, where an auctioneer (*i.e.*, spectrum authority) sells licenses for signal transmission on available channels to n bidders (*i.e.*, spectrum users) located in a geographic region. We assume channels over the spectrum to be auctioned have uniform characteristics and values, and each bidder i desires d_i channels. In practice, channels are auctioned for use over a period of time, and they will be dynamically occupied and vacated by winning bidders. Due to the dynamic nature of channel occupancy and release, the total channel supply is uncertain. A key characteristic of wireless communications is the ability to re-use frequencies/channels to increase system capacity, so different bidders sufficiently far apart can operate on the same frequency/channel. To model this, we represent the interference condition among bidders by an interference graph: two bidders either interfere with each other or can use the same channel simultaneously. We assume bidders do not collude with each other and make their bids independently. When idle channels are arriving dynamically, the auctioneer should make the allocation and payment decisions in an online manner, with the goal of achieving price fairness, truthfulness and bounded performance.

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

Channel demand (d_i) – It represents the number of channels bidder i would like to bid. Let $\mathcal{D} = \{d_1, d_2, ..., d_n\}$ denote the set of channel demands of all the bidders.

Per-channel valuation (v_i) – It represents the true price bidder i is willing to pay for one channel. We assume that the per-channel valuation is a private value and is known only to the bidder itself.

Per-channel payment (p_i) – It represents the bidder i's payment for one channel.

Bidder utility (u_i) – The *utility* of bidder i is defined by $u_i = v_i \cdot x_i - p_i \cdot x_i$, where x_i denotes the number of channels bidder i gets after the spectrum auction.

Definition 1. *A truthful spectrum auction is the one that for any spectrum user i, regardless of the declarations of the others, u_i achieves the maximum when user i bids for each channel at its valuation, i.e., $b_i = v_i$.*

To achieve truthfulness, many spectrum auction designs determine each winner's payment based on other users' bids. As a consequence, the same channel is most probably priced at different values for different winning users. From the perspective of users, there exists price discrimination for *identical goods* (*i.e.*, homogeneous channels) such that a winner charged at a higher price may prefer another winner's allocated channels and payment to his own. To this end, we have the following definition.

Definition 2. *A fair spectrum pricing scheme is the one that each winning spectrum user will pay exactly or almost the same price for identical channels.*

One of our design goals is to maximize the *social welfare*, which is the sum of the valuations of winning bidders. In our model, due to spectrum reusability, a channel simultaneously used by two interfere-free users can be considered as two distinct selling items. For ease of understanding and explanation, in the following discussion we assume the number of selling items is l when taking into account the effect of reusability. Then, the *optimal social welfare* should be carefully defined to characterize our model. In the following, we first give the definition of optimal social welfare for the single-unit demand case. Without loss of generality, we assume that $v_1, v_2, \ldots, v_n$ are sorted in non-increasing order.

Definition 3. *The optimal social welfare for the single-unit demand case is denoted by $S\text{-}OPT_L^{(l)} = \sum_{i=1}^{l} v_i$ when l selling items could be provided to L winners.*

In single-unit case, it is easy to see that $L = l$ when taking into account reusability. So, we use $S\text{-}OPT_L$ instead of $S\text{-}OPT_L^{(l)}$ for simplicity. We next consider the optimal solution for the multi-unit demand case. Unlike the single-unit case, the last winner may be partially-satisfied due to $\sum_{i=1}^{L} d_i > l$. If we consider a strictly single-minded user (*i.e.*, this partially-satisfied bidder will not be regarded as a winner), given l selling items, the social welfare maximization problem will become a knapsack problem [12]. Consequently, this will pose a challenge for defining the optimality in our model since the optimal solution of knapsack problem conflicts with the fairness of allocation, *e.g.*, the partially-satisfied user with higher valuation may be replaced by users with lower valuations. Therefore, we weaken the strict demand requirement and allow that the last winner could be partially-satisfied. Then we define the optimal solution as follows.

Definition 4. *The optimal social welfare for the multi-unit case is denoted by* $M\text{-}OPT_L^{(l)} = \sum_{i=1}^{L-1} v_i d_i + v_L(l - \sum_{i=1}^{L-1} d_i)$, *subject to* $L = \arg\min_{1 \le L \le n}(\sum_{i=1}^{L} d_i \ge l)$ *when* l *items could be provided to* L *winners.*

In the above definition, $(l - \sum_{i=1}^{L-1} d_i)$ denotes the selling items obtained by the last winner, $v_L(l - \sum_{i=1}^{L-1} d_i)$ denotes the total valuation of the last winner, and the sum of valuations of the $L-1$ fully-satisfied winners is denoted by $\sum_{i=1}^{L-1} v_i d_i$. Note that the last winner could either be fully-satisfied or partially-satisfied.

III. Our Spectrum Auction Construction

In this section, we first illustrate the deficiencies and challenges in designing efficient and truthful spectrum auctions supporting price fairness under dynamic spectrum supply. Then, we propose new and novel spectrum auction constructions achieving all desirable properties, including truthfulness, price fairness and online auction with dynamic supply.

A. Identifying the Challenges of Spectrum Auction Design

We show that the existing truthful spectrum auction designs are insufficient to meet the above properties when applied to our model. For the ease of understanding, we assume each bidder requests at most one channel. The same conclusions could be applied to the multi-unit case.

1) The Insufficiency of Price Fairness with A Dynamic Spectrum Supply: Consider the classical k-item Vickrey auction [13] where k winners pay at the $(k+1)$-th bidder's bid, it achieves price fairness according to Definition 2. We show that when it is applied to a dynamic supply setting, the pricing outcomes become unfair. Assume there are 3 bidders, each of which requests one channel and interferes with each other. In a dynamic supply setting, there are 2 idle channels arriving sequentially. Without loss of generality, we assume $\mathcal{B} = \{b_1 = 7, b_2 = 6, b_3 = 1\}$. Because idle channels should be instantly allocated in an online manner, the first idle channel will be allocated to bidder one with the pricing value equivalent to bidder two's bid, *i.e.*, $p_1 = b_2 = 6$. After a short time period, the second channel is available and will be allocated to bidder two, and the payment equals to bidder three' bid, *i.e.*, $p_2 = b_3 = 1$. During one auction period, the pricing is obviously unfair since bidder one values the channel most (*i.e.*, declares the highest bid) but pays much more than bidder two for an idle channel. Thus, the natural extension of Vickrey auction to the online case with a spectrum dynamic supply makes the pricing unfair.

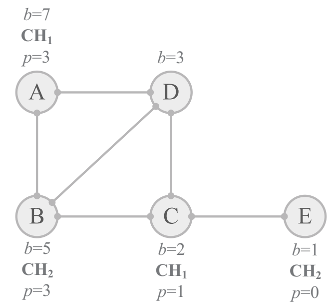

Fig. 1. An example of price unfairness when idle channels arrive at the same time point.

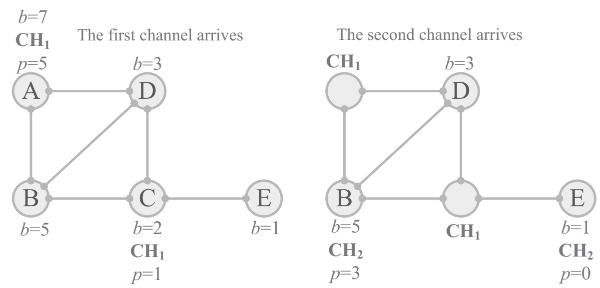

Fig. 2. An example of price unfairness when idle channels arrive at different time points.

We next consider another truthful and efficient auction mechanism named VERITAS, which is particularly designed for spectrum auctions with channel reusability [1].

Allocation:

- Sort the bids in non-increasing order.
- Extract the first bidder (*e.g.*, bidder i) in the sorted list and check whether there exists a channel to satisfy bidder i, *i.e.*, $|Distinct(N(i))| + 1 \le k$, where $|Distinct(N(i))|$ denotes the number of channels allocated to the interfering bidders of bidder i, and k is the number of independent channels. If the checking equation holds, allocate bidder i a channel with the lowest available index not in $Distinct(N(i))$.
- Repeat step 2 until all the bidders are examined.

Pricing:

- Find the critical neighbor for each winner. The critical neighbor of bidder i is defined as follows: bidder i can get allocated if and only if it bids higher than its critical neighbor.
- Charge each winner i with the bid of its *critical neighbor* multiplied by the number of channels allocated to winner i.

In Figs. 1 and 2, we show that VERITAS cannot achieve price fairness using two counter examples. Assume there are 5 bidders (A, B, C, D, E) with bids $\{b_A = 7, b_B = 5, b_C = 2, b_D = 3, b_E = 1\}$, each requesting at most one channel.

Figs. 1 and 2 show the conflict graphs of 5 bidders competing for two channels CH_1 and CH_2. In Fig. 1, two channels arrive at the same time point. After the greedy allocation and the critical value based pricing, bidders A, B, C and E get allocated. However, for homogenous channels the payment of the winners are quite different: bidders A and B pay more than twice as the payment of bidder C, and more seriously bidder E pays nothing for free use. In Fig. 2, two channels arrive sequentially at different times. Compared to the case in Fig. 1, while the winners are still A, B, C and E, bidder A pays even more and is charged at 5. Obviously, for both scenarios VERITAS [1] contradicts the definition of price fairness.

2) Truthfulness Under Channel Reusability: When directly applied to spectrum auctions, conventional truthful auction designs, such as secondary pricing spectrum auction and VCG-style spectrum auction, become untruthful due to channel reusability [1]. Unlike conventional auction models, spectrum bidders have interference constraints with each other such that the number of available channels is different for different bidders. Therefore, a bidder could manipulate its bid to disrupt the resource allocation and pricing and thus inherently violates the truthfulness of auction designs. Keeping fairness and efficiency in mind, we should also carefully take care of truthfulness under channel reusability.

All of the above observations highly motivate us to design a new yet practical spectrum auction system that achieves all desirable properties.

B. Spectrum Auction with Dynamic Supply: The Single-unit Case

In this section, we present a new spectrum auction scheme, achieving truthfulness, fairness with dynamic supply and efficiency. We start from the single-unit case, where each spectrum user requests at most one channel. In the next subsection, we will extend our approach to support multi-unit demand users.

Different from the existing spectrum auction schemes, our scheme will choose some bidders as *eligible bidders* who receive equal opportunity to get allocated. Intuitively, the larger value a user bids at, the higher chance it becomes an eligible bidder. However, an eligible user may lose a bid due to the introduction of randomization in the allocation process.

Our spectrum auction scheme mainly includes three parts: an initialization process, a spectrum winner selection algorithm and a pricing algorithm. Our spectrum auction design for the single-unit demand case is shown in Algorithm 1. We denote bidder i's bid by $[b_i, d_i]$, where b_i is the per-channel bid and d_i is the number of channels requested by i. In the single-unit case, $d_i = 1$. Assume that the bid set $\mathcal{B}$ is sorted in non-increasing order of b_i. Without loss of generality, let $\mathcal{B} = \{b_1, b_2, \ldots, b_n\}$, where n denotes the total number of bidders. The initialization mainly consists of two steps:

1) *Eligible bidder selection.* Select q from $\{2^1, 2^2, \ldots, 2^i, \ldots, n\}$ uniformly at random and let the q top-ranking bidders in $\mathcal{B}$ be *eligible bidders*.
2) *Bidder grouping.* Divide eligible bidders into multiple interference-free groups $\mathcal{G} = \{g_1, g_2, \ldots, g_m\}$ using graph coloring algorithm [14].

By selecting the q top-ranking bidders in $\mathcal{B}$ as eligible bidders, no user will be an eligible bidder before others whose per-channel bids are higher than its per-channel bid anytime. This means that the larger the bid, the higher probability the bidder will be selected as an eligible bidder. In other words, the selection of eligible bidders is to assign idle channels to the top-ranking bidders in an online manner. Note that, because grouping first before eligible bidder selection will generally result in more bidder groups (the grouping approach is independent of bid values), ineligible bidders that interference with eligible bidders will decrease the spectrum utilization. This motivates us to select eligible bidders first before performing the bidder grouping.

Before presenting the *winner selection* step, we first give an important definition used in our algorithm.

Definition 5. *An* ***unassigned group*** *is a group in which none of bidders has got allocated. An* ***unassigned group*** *becomes an* ***assigned group*** *once a bidder of the group gets allocated an idle channel, and accordingly the channel allocated to the bidder is assigned to this group.*

Note that, even if a bidder belongs to an *assigned group*, it will not get allocated the channel assigned to its group until it is activated first. In our spectrum auction model, idle channels are arriving in an unpredictable manner, the spectrum auctioneer should choose winners and allocate channels online:

1) Obtain a random permutation of $\{b_1, b_2, \ldots, b_q\}$, denoted by $\mathcal{B}' = \{b'_1, b'_2, \ldots, b'_q\}$.
2) Extract the first available bidder in $\mathcal{B}'$ and check if the group to which it belongs is *unassigned* or not. If the group is an *unassigned* one, allocate directly an idle channel to it and mark the group as *assigned*; otherwise activate the bidder and allocate it the channel that has been already assigned to its group. Eliminate this satisfied bidder from $\mathcal{B}'$.
3) Continue step 2 until all idle channels (which arrive online) are allocated or all q eligible bidders have been satisfied.

In Algorithm 1, the motivation of permutating per-channel bids of eligible bidders is that, eligible bidders may have the motivation to improve their own bids to obtain spectrum resource as early as possible due to the dynamical arrival of idle channels. Another key point in our design is that an eligible bidder will not get allocated until being activated. This is to ensure the channel allocation sequence strictly follows the permutation order of spectrum bidders.

We next consider the pricing algorithm. In online auctions with a dynamic spectrum supply, when bidder i gets allocated, it should be charged at the same time. The random number q obtained during *eligible bidder selection* determines not only the maximum number of winners but also the prices to be charged for all winners. So, no matter which group the bidder belongs to, each winning bidder is charged at b_{q+1}, which is equivalent to the highest bid declined by the bidder (which is not allowed to use idle spectrum resources):

$$p_i = b_{q+1}.$$

Bidders that do not get allocated will pay zero.

Algorithm 1 Online spectrum auction for the single-unit demand user under an unknown supply of idle channels.

Unassigned group set: $UG = \{g_1, g_2, \ldots, g_m\}$;
Assigned group set: $AG = \varnothing$
Sort bids in $\mathcal{B}$ and obtain $\{b_1, b_2, \ldots, b_n\}$;
Select q randomly from $\{2^1, 2^2, \ldots, 2^i, \ldots, n\}$;
Divide $\{b_1, \ldots, b_q\}$ into interference-free bidder groups;
Permutate $\{b_1, b_2, \ldots, b_q\}$ to get $\mathcal{B}'$; *%initialization is done, winner selection and channel allocation starts as follows*
for each idle channel j coming dynamically **do**
  $avail\_ch = \{c_j\}$;
  if $\mathcal{B}' = \varnothing$ **then**
    break;
  else
    $b = FirstBid(\mathcal{B}')$; *%get the first bid in* $\mathcal{B}'$
    $i = ID(b)$; *%get ID of the bidder with bid* $b^{(i)}$
    $g^{(i)} = Group(i)$; *%find the group which bidder i belongs to*
    while $(avail\_ch \neq \varnothing)$ **or** $(g^{(i)} \in AG)$ **do**
      if $g^{(i)} \in UG$ **then**
        $Allocate(i)$;
        $p_i = b_{q+1}$; $\%b_{q+1} = 0$ *when* $q = n$
        $Assign(c_j, g^{(i)})$;
        $avail\_ch = avail\_ch \setminus \{c_j\}$;
        $UG = UG \setminus \{g^{(i)}\}$;
        $AG = AG \cup \{g^{(i)}\}$;
      else
        $Activate(i)$ and $Allocate(i)$;
        $p_i = b_{q+1}$;
      end if
      $B' = B' \setminus \{b\}$;
      $b = FirstBid(\mathcal{B}')$;
      $i = ID(b)$;
      $g^{(i)} = Group(i)$;
    end while
  end if
end for

We next analyze the properties of the proposed spectrum auction scheme dealing with the single-unit demand case.

Analysis of truthfulness: In the following, we prove the truthfulness by showing whether or not all eligible bidders can be satisfied, a bidder cannot gain any benefits by untruthful bidding.

Lemma 1. *When all eligible bidders can be satisfied, any bidder cannot misreport the per-channel bid to increase its utility.*

Proof: Assume there exist enough idle channels arriving dynamically such that all eligible bidders will be satisfied during one auction period. Among n bidders, however, only some of them will become eligible bidders. Let u_i^v and u_i^b be bidder i's utilities when bidding at v_i and b_i respectively, where v_i denotes bidder i's valuation. By the definition of truthfulness, we should ensure $u_i^b \leq u_i^v$ for $b_i \neq v_i$ in all cases. In our auction scheme, $p_i = b_{q+1}$ and q is chosen randomly. Then, we show the truthfulness of our auction scheme by analyzing all possible cases.

- Case 1: $b_i > v_i > b_{q+1}$. No matter whether bidder i bids at b_i or v_i, it will always be an eligible bidder, and the clearing price b_{q+1} will be the same. Thus, $u_i^b = u_i^v = v_i - b_{q+1}$, the claim holds.
- Case 2: $b_i > v_i = b_{q+1}$. No matter whether bidder i bids at b_i or v_i (if bidding at v_i, it will not be selected as an eligible bidder if there are strictly q bidders with bids larger than b_{q+1}; however, if among the top-q ranking bids, some of them have bids that equal to b_{q+1}, it may be selected as an eligible bidder), the utility of bidder i will always be zero, $u_i^b = v_i - b_{q+1} = 0 = u_i^v$, the claim holds.
- Case 3: $v_i > b_i > b_{q+1}$. No matter whether bidder i bids at b_i or v_i, it will definitely be satisfied. Then, the utility of bidder i will be $u_i^b = u_i^v = v_i - b_{q+1}$.
- Case 4: $v_i > b_i = b_{q+1}$. Similar to Case 2, when bidder i bids at b_i, its utility can be 0 or $u_i^b = v_i - b_{q+1}$. However, when bidding at v_i, it will be an eligible bidder for sure with $u_i^v = v_i - b_{q+1}$. So we can claim that $u_i^v \geq u_i^b$.
- Case 5: $v_i > b_{q+1} > b_i$. Because $b_i < b_{q+1}$, bidder i will not be an eligible bidder when bidding at b_i and thus $u_i^b = 0$. However, bidder i will be satisfied when bidding at v_i. Hence, $u_i^b < u_i^u = v_i - b_{q+1}$.
- Case 6: $b_i > b_{q+1} > v_i$. When bidder i bids truthfully, it will not be an eligible bidder and thus $u_i^v = 0$. It will be satisfied when misreporting the bid, but $u_i^b = v_i - b_{q+1} < 0$. Hence, $u_i^b < u_i^v$.
- Case 7: $b_{q+1} = b_i > v_i$. When bidding at b_i, if bidder i is selected as an eligible bidder, the utility $u_i^b = v_i - b_{q+1} < 0$. When bidding at v_i, bidder i will not be eligible bidder and thus $u_i^v = 0$.
- Case 8: $b_{q+1} > b_i > v_i$. Bidder i will not be satisfied no matter it bids at b_i or v_i, so $u_i^b = u_i^v = 0$.
- Case 9: $b_{q+1} \geq v_i > b_i$. The subcases of this case are similar to Cases 7 and 8.

In summary, if all eligible bidders can be satisfied by the incoming idle channels, we show that a bidder achieves maximum utility when bidding truthfully. This completes the proof. ■

We next consider the case where only some of the eligible bidders can be satisfied due to the limited number of idle channels during one auction time period. In the *winner selection and channel allocation* phase, if a greedy allocation method (based on the per-channel bids) is adopted, an eligible bidder is highly motivated to improve its bid to obtain a higher rank so as to get allocated earlier in the eligible bidder set. Therefore, the order of allocation must be bid-independent.

Lemma 2. *When not all eligible bidders can be satisfied, a bidder cannot misreport the per-channel bid to increase its utility.*

Proof: Assume there does not exist a sufficient number of idle channels that can satisfy all eligible bidders during the online allocation. Thus, all eligible bidders have the motivation to get allocated earlier to prevent it encountering resource deficiency, which will lead to zero utility. In our algorithm, we randomly permutate the sorted list of eligible bidders to disrupt the order of allocation among them such that the order of allocation is independent of the bid values of all eligible bidders. Hence, an eligible bidder cannot overbid to increase the probability of getting allocated earlier, *i.e.*, overbidding will not help to increase its utility. On the other hand, by underbidding, an eligible bidder may lose the opportunity to get allocated and its utility will be zero.

As shown in Lemma 1, even if the channel supply is sufficient, an ineligible bidder cannot increase its utility by misreporting bid. Thus, we omit the discussions with respect to ineligible bidders here.

■

Theorem 3. *Under the dynamic channel supply, the proposed spectrum auction for the single-unit case is truthful.*

Proof: By combining Lemmas 1 and 2, we can conclude that the proposed spectrum auction scheme for the single-unit case is truthful. ■

Analysis of fairness: The following theorem shows the fairness of our spectrum auction scheme.

Theorem 4. *Under the dynamic channel supply, the proposed spectrum auction achieves price fairness.*

Proof: In our algorithm, each winner will pay b_{q+1} for one channel it gets, *i.e.*, the per-channel payments are the same for all winners. According to Definition 2, our spectrum auction scheme for the single-unit case achieves fair pricing. ■

Analysis of optimality: The following theorem characterizes the approximation ratio.

Theorem 5. *Under dynamic channel supply, the proposed spectrum auction for the single-unit case can achieve a $\log n$ approximation to the optimal social welfare.*

Proof: We first consider the optimal allocation that achieves the maximized social welfare. Let $\mathbb{P}$ denote the set of all possible random permutations for all bidders, P_k denote the k_{th} permutation. Because the number of selling items is determined by the allocation results, we use R_k to denote the number of selling items for the k_{th} permutation. Let $R_{\min} = \min\{R_k | P_k \in \mathbb{P}\}$ be the minimum number of selling items among all possible permutations. Since the *winner selection* (*i.e.*, random permutation of eligible bidders) is a randomized scheme and the number of selling items is uncertain, by using $R_{\min}$, we can guarantee that $R_{\min}$ selling items could be provided for any permutation or say resulting allocation. We then adopt $R_{\min}$ in our analysis to evaluate the optimal solution by taking all possible permutations and interference constraints between bidders into account. We denote the social welfare of the $R_{\min}$ highest bidders by $S\text{-}OPT_{R_{\min}} = \sum_{i=1}^{R_{\min}} v_i$. In *eligible bidder division*, q is chosen randomly from $\{2^1, 2^2, \ldots, 2^i, \ldots, n\}$ to divide all bidders into eligible bidders and ineligible bidders. It is easy to see that the probability of choosing each possible value of q is $1/\log n$. In the following, we show that it is sufficient to analyze two special cases to derive the approximation ratio.

- Case 1: $R_{\min} < q \le 2R_{\min}$. In this case, at least half of the q bidders will be selected as winners. Because $R_{\min}$ winners are randomly chosen among eligible bidders, the expectation of the social welfare is $\frac{R_{\min}}{q} \cdot S\text{-}OPT_q$. When $R_{\min} < q \le 2R_{\min}$, the social welfare is equal to or greater than $\frac{1}{2} S\text{-}OPT_q$.
- Case 2: $\frac{1}{2}R_{\min} < q \le R_{\min}$. In this case, there exists a surplus of selling items to be allocated to eligible bidders, so all eligible bidders will be satisfied. The social welfare thus is $S\text{-}OPT_q$.

We denote the social welfares of the case 1 and the case 2 by SW_1 and SW_2, respectively. In case 1, because $q > R_{\min}$, we have $SW_1 \ge \frac{1}{2} S\text{-}OPT_q > \frac{1}{2} S\text{-}OPT_{R_{\min}}$. In case 2, because $q > \frac{1}{2}R_{\min}$ and v_is are sorted in a non-increasing order, we have $SW_2 = S\text{-}OPT_q > \frac{1}{2} S\text{-}OPT_{R_{\min}}$. Thus, the social welfare of the proposed spectrum auction scheme *in expectation* is lower-bounded by

$$
\begin{aligned}
\mathbf{E}[SW] &= (1/\log n) \cdot (\sum_{i=1}^{\log n} SW_i) \qquad (1)\\
&\ge (1/\log n) \cdot (SW_1 + SW_2)\\
&> (1/\log n) \cdot (\frac{1}{2} S\text{-}OPT_{R_{\min}} + \frac{1}{2} S\text{-}OPT_{R_{\min}})\\
&= (1/\log n) \cdot S\text{-}OPT_{R_{\min}}
\end{aligned}
$$

■

C. Spectrum Auction with Dynamic Supply: The Multi-unit Case

In this section, we discuss the multi-unit case where each bidder may demand more than one channels. In the following, we will show that the spectrum auction design for multi-unit case requires us to tackle many unique challenges.

Similar to the single-unit case, during *initialization*, we first sort all bidder requests by their per-channel bids in a non-increasing order $\mathcal{B} = \{[b_1, d_1], [b_2, d_2], \ldots, [b_n, d_n]\}$, where d_i denotes the channel demand of bidder i. However, we will represent each *original bidder* by a set of *virtual bidders*, *i.e.*, a bidder who demands d_i channels will be convert into d_i virtual bidders, each of which demands only one channel. Each *virtual bidder* inherits the interference constraints of its own father, *i.e.*, the *original bidder*, and conflicts with other virtual bidders generated from the same original bidder. Obviously, the *virtual request* can be represented as $\mathcal{B}_v = \{b_1^1, \ldots, b_1^{d_1}, b_2^1, \ldots, b_2^{d_2}, \ldots, b_n^1, \ldots, b_n^{d_n}\}$, and the number of virtual bidders is $m = \sum_{i=1}^{n} d_i$. Follow the auction design of the single-unit case, we select a random q uniformly from $\{2^1, 2^2, \ldots, 2^i, \ldots, m\}$, where the q top-ranking virtual bidders in $\mathcal{B}_v$ will be *eligible virtual bidders*. Due to the construction of $\mathcal{B}_v$, it is easy to see that the top-ranked q virtual bidders are corresponding to a set of top-ranked original bidders in $\mathcal{B}$. Then, we perform the *bidder grouping* based on eligible virtual bidders.

After *initialization*, the *winner selection* is executed as follows. Firstly, to ensure a winning bidder's multi-unit demand can be satisfied, we perform the random permutation over *original bidders* instead of *virtual bidders*, generating $\mathcal{B}'$. When idle channels arrive, we propose to satisfy the eligible original bidders in sequence according to $\mathcal{B}'$. Specifically, we allocate channel supply to virtual bidders of an eligible original bidder, and an eligible original bidder can be satisfied when its corresponding virtual bidders are satisfied. Then, we choose the next eligible original bidder for channel allocation. In the *winner selection* and *channel allocation* process, like the operations in the single-unit case, we still need to check if the group a virtual bidder belongs to is an *assigned* one or an *unassigned* one, based on which we activate the virtual bidder before allocating it an idle channel. Due to the space limitation, we will not go into details.

We next discuss the pricing of multi-unit demand spectrum users. We show that if we follow the same pricing strategy by charging all winning virtual bidders at b_{q+1}, spectrum users may have the motivation to improve their utility by misreporting their bids.

An illustrating example. Assume there are three bidders A, B and C, with channel demands $d_A = 3$, $d_B = 2$, and $d_C = 1$, respectively. Their per-channel valuations or say true per-channel bids are $v_A = 3$, $v_B = 2$ and $v_C = 1$. If the number of eligible virtual bidders is chosen as $q = 4$, all the three virtual bidders of A (A_1, A_2, A_3) and one virtual bidder of B (B_1) are selected as eligible virtual bidders. Following the single-unit case, all the winners will pay at the bid of the $q + 1_{\text{th}}$ virtual bidder. Thus, $p_{A_1} = p_{A_2} = p_{A_3} = b_{B_2} = 2$, and the utility of A is

$$u_A = 3 \cdot v_A - 3 \cdot b_{B_2} = 9 - 3 \cdot 2 = 3.$$

However, if bidder A underbids at 1.1 for each channel, the $q + 1_{\text{th}}$ virtual bidder falls into the virtual bidders of A. The resulting utility of bidder A is

$$u_A = 2 \cdot v_{A_1} - 2 \cdot b_{A_3} = 6 - 2 \cdot 1.1 = 3.8.$$

As can be seen, bidder A is partially satisfied and obtains a higher utility by manipulating its bid. In the following, we determine the *per-channel price* for each eligible original bidder and show that their prices are *almost* the same. We first give two important notations.

$x_j(q)$ – Given a randomly-chosen q, it represents the number of channels a winning original bidder j can get at most.

$x_j^{(-i)}(q)$ – Given a randomly-chosen q, it represents the number of channels a winning original bidder j can get at most when bidder i does not participate the auction.

For an ineligible or say a losing bidder j, we define $x_j(q) = x_j^{(-i)}(q) = 0$. Then, before *channel allocation*, we can compute the per-channel price for each winning original bidder in advance as

$$p_i = \frac{\sum_{j \neq i}^{n} x_j^{(-i)}(q) \cdot b_j - \sum_{j \neq i}^{n} x_j(q) \cdot b_j}{x_i(q)}. \qquad (2)$$

Next, we analyze the properties of our spectrum auction scheme for multi-unit demand case.

Analysis of fairness: Different from the single-unit case where all winners pay *exactly* the same (per-channel) price for homogeneous channels, in the multi-unit case we show the per-channel prices for all winners are *almost* the same for homogeneous channels. Based on Eq. (2), we can derive the total payment of bidder i as

$$x_i(q) \cdot p_i = \sum_{j \neq i}^{n} x_j^{(-i)}(q) \cdot b_j - \sum_{j \neq i}^{n} x_j(q) \cdot b_j. \qquad (3)$$

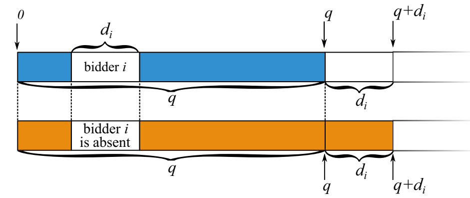

Fig. 3. The illustration of a bidder's payment.

Fig. 3 shows the distribution of channel demands when bidder i does (in blue) and does not (in orange) participate the auction. In conjunction with Eq. (3), it can be seen that the total payment of bidder i (*i.e.*, $x_i(q) \cdot p_i$) is the social welfare generated from bidders in the "orange" area $[q, q + d_i]$. Thus, the per-channel price for bidder i can be considered as the average of bids of bidders located in this area, *i.e.*, the ratio of the social welfare generated from bidders in the "orange" area to d_i.

In our auction design, bidders in the "orange" area $[q, q + d_i]$ are sorted by their per-channel bids. When per-channel bids of spectrum users are distributed uniformly at random, for bidders i and j with demands $d_i \neq d_j$, the *average bids* of bidders located in the d_i-length orange area and the d_j-length orange area are pretty close to each other. Assume there are 100 bidders, whose per-channel bids are randomly distributed over $(0, 1]$ and demands are randomly chosen from $\{1, 2, 3\}$. We select a random q from $\{2^1, 2^2, \ldots, 2^i, \ldots, m\}$, where m is the total number of demands of all bidders. Fig. 4 shows the per-channel price for all winners are almost the same.

Analysis of truthfulness: In the following, we first prove our auction design has individual rationality, based on which we show eligible bidders (whether satisfied or not) cannot gain any benefits by untruthful bidding.

Lemma 6. *Our spectrum auction scheme for the multi-unit case achieves individual rationality, i.e., no winning original bidder pays more than its total bid for the allocated channels:* $x_i(q) \cdot b_i \geq x_i(q) \cdot p_i$ *or* $u_i \geq 0$.

Proof: If bidder i does not get allocated, then $x_i(q) = 0$, the lemma holds. Otherwise, since we sort bidders by their per-channel bids, based on which we select eligible virtual bidders, it is easy to see the social welfare $\sum_{i=1}^{n} x_i(q) \cdot b_i \geq \sum_{j \neq i}^{n} x_j^{(-i)}(q) \cdot b_j$. Thus, we have

$$\sum_{i=1}^{n} x_i(q) \cdot b_i \geq \sum_{j \neq i}^{n} x_j^{(-i)}(q) \cdot b_j$$

$$\sum_{j \neq i}^{n} x_j(q) \cdot b_j + x_i(q) \cdot b_i \geq \sum_{j \neq i}^{n} x_j^{(-i)}(q) \cdot b_j$$

$$x_i(q) \cdot b_i \geq \sum_{j \neq i}^{n} x_j^{(-i)}(q) \cdot b_j - \sum_{j \neq i}^{n} x_j(q) \cdot b_j$$

$$x_i(q) \cdot b_i \geq x_i(q) \cdot p_i \tag{4}$$

Hence, Lemma 6 holds.

■

We next prove that our spectrum auction scheme for multi-unit demand users is truthful, *i.e.*, spectrum users cannot increase it's utility by bidding untruthfully. Due to Eq. (2), we can compute the utility of bidder i as

$$\begin{aligned} u_i &= x_i(q) \cdot v_i - x_i(q) \cdot p_i \qquad (5) \\ &= x_i(q) \cdot v_i - \sum_{j \neq i}^{n} x_j^{(-i)}(q) \cdot b_j + \sum_{j \neq i}^{n} x_j(q) \cdot b_j. \end{aligned}$$

In Eq. (5), because the second item $\sum_{j \neq i}^{n} x_j^{(-i)}(q) \cdot b_j$ does not depend on i (it denotes the social welfare where bidder i does not participate the auction), we could ignore it in the analysis below. Thus, we have

$$u_i \sim v_i \cdot x_i(q) + \sum_{j \neq i}^{n} x_j(q) \cdot b_j, \tag{6}$$

where $\sim$ means u_i only relates to the right two terms.

In the following, we first analyze the case where there exists a sufficient number of idle channels such that all eligible (original) bidders can be satisfied.

Lemma 7. *A fully-satisfied eligible bidder cannot increase its utility by manipulating its bid.*

Proof: A fully-satisfied eligible original bidder i gets d_i channels it demands. If it manipulates its bid, there will be three possible cases: i) Bidder i remains to be a fully-satisfied eligible bidder, then $x'_i(q) = x_i(q)$. According to our *winner selection* process, it will not affect $x_j(q) (j \neq i)$ of any other winning bidder j. So, $\sum_{j \neq i}^{n} x_j(q) \cdot b_j$ remains the same. Then, according to Eq. (6), we have $u'_i = u_i$; ii) Bidder i becomes a partially-satisfied eligible bidder (by underbidding), then $x'_i(q) < x_i(q)$ and $\sum_{j \neq i}^{n} x'_j(q) \cdot b_j > \sum_{j \neq i}^{n} x_j(q) \cdot b_j$. Given q, this implies some ineligible bidders get the opportunity to join the eligible bidder set. Obviously, the per-channel bids of these bidders are smaller than the per-channel valuation of bidder i. Thus, we have

$$\sum_{j \neq i}^{n} x'_j(q) \cdot b_j - \sum_{j \neq i}^{n} x_j(q) \cdot b_j \leq v_i \cdot x_i(q) - v_i \cdot x'_i(q)$$

$$v_i \cdot x'_i(q) + \sum_{j \neq i}^{n} x'_j(q) \cdot b_j \leq v_i \cdot x_i(q) + \sum_{j \neq i}^{n} x_j(q) \cdot b_j.$$

Based on Eq. (6), the above inequality can be expressed as $u'_i \leq u_i$; iii) Bidder i becomes an ineligible bidder, the utility $u'_i = 0$. According to Lemma 6, $u_i \geq 0$. Thus, we have $u_i \geq u'_i$.

■

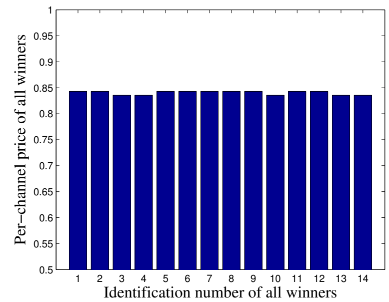

Fig. 4. The comparison of per-channel payments of all winners.

Lemma 8. *A partially-satisfied eligible original bidder cannot increase its utility by manipulating its bid.*

Proof: In our algorithm, since q is randomly chosen and pre-determined before the start of spectrum allocation, there may exist one partially-satisfied eligible original bidder i. If it manipulates its bid, there will be three possible cases: i) Bidder i remains to be a partially-satisfied eligible bidder. then $x'_i(q) = x_i(q)$. Similarly to the proof of case 1 in Lemma 7, we have $u'_i = u_i$; ii) Bidder i becomes a fully-satisfied eligible bidder (by overbidding), then we have $x'_i(q) > x_i(q)$ and $\sum_{j \neq i}^{n} x'_j(q) \cdot b_j < \sum_{j \neq i}^{n} x_j(q) \cdot b_j$. The difference between $\sum_{j \neq i}^{n} x_j(q) \cdot b_j$ and $\sum_{j \neq i}^{n} x'_j(q) \cdot b_j$ is the total bids of virtual bidders who are preempted by virtual bidders of bidder i. Given q, the number of preempted channels is equal to the number of channels bidder i additionally obtain by preemption. Obviously, the per-channel bids of bidders being preempted are larger than the per-channel valuation of bidder i. Thus, we have

$$\sum_{j \neq i}^{n} x_j(q) \cdot b_j - \sum_{j \neq i}^{n} x'_j(q) \cdot b_j \geq v_i \cdot x'_i(q) - v_i \cdot x_i(q)$$

$$v_i \cdot x_i(q) + \sum_{j \neq i}^{n} x_j(q) \cdot b_j \geq v_i \cdot x'_i(q) + \sum_{j \neq i}^{n} x'_j(q) \cdot b_j.$$

Based on Eq. (6), the above inequality can be expressed as $u_i \geq u'_i$; iii) Bidder i becomes an ineligible bidder, the utility $u'_i = 0$. According to Lemma 6, $u_i \geq 0$. Thus, we have $u_i \geq u'_i$. ■

Lemma 9. *An ineligible bidder cannot increase its utility by manipulating its bid.*

Proof: An ineligible bidder i will not get allocated, then we have $x_i(q) = 0$ and $u_i = 0$. If an ineligible bidder untruthfully bids (by overbidding) to join the eligible bidder set, some eligible bidders will become ineligible and lose their bids. Obviously, the average per-channel bid of these losing bidders is greater than or equal to the per-channel valuation of bidder i (otherwise they would be ineligible bidders when bidder i bids truthfully). Thus, we have

$$\frac{\sum_{j\neq i}^{n} x_j(q)\cdot b_j - \sum_{j\neq i}^{n} x'_j(q)\cdot b_j}{x'_i(q)} \geq v_i \tag{7}$$

Based on Eq. (6), we can compute the utility difference

$$\begin{aligned} u'_i - u_i &= v_i \cdot x'_i(q) + \sum_{j\neq i}^{n} x'_j(q)\cdot b_j - \\ &\quad (v_i\cdot x_i(q) + \sum_{j\neq i}^{n} x_j(q)\cdot b_j). \end{aligned}$$

Because $x_i(q) = 0$ and $u_i = 0$, we have

$$u'_i = v_i\cdot x'_i(q) + \sum_{j\neq i}^{n} x'_j(q)\cdot b_j - \sum_{j\neq i}^{n} x_j(q)\cdot b_j.$$

Based on Eq. (7), we have $u'_i \leq 0$. ■

We next analyze the case where not all eligible (original) bidders can be satisfied.

Lemma 10. *When not all eligible original bidders can be satisfied, an original bidder cannot misreport the per-channel bid to increase its utility.*

Proof: The proof is similar to Lemma 2, and thus we omit the details of the proof here. ■

Theorem 11. *Under the dynamic channel supply, the proposed spectrum auction for the multi-unit case is truthful.*

Proof: By combining Lemmas 7, 8, 9 and 10, it can be see that any bidder has no incentive to misreport their bids to improve its utility. According to Definition 1, our spectrum auction scheme for the multi-unit demand case is truthful. ■

Analysis of optimality: The following theorem characterizes the approximation ratio.

Theorem 12. *Under the dynamic channel supply, the proposed spectrum auction for the multi-unit case can achieve $(\frac{T}{Q+1} + \frac{1}{2})\log m$ approximation to the optimal social welfare.*

Proof: Similar to the single-unit case, we first consider the optimal allocation that achieves the maximum social welfare. To obtain the optimal allocation, we assume that all original bidders are eligible to be allocated and there exist $R'_{\min}$ selling items such that $R'_{\min}$ virtual bidders will be satisfied. Similar to the proof of Theorem 5, we use $R'_{\min} = \min\{R'_k|P_k \in \mathbb{P}\}$ to denote the minimum number of selling items for all possible permutations (allocations).

Assume the $R'_{\min}$ virtual bidders are corresponding to $R_{\min}$ original bidders, we denote the optimal social welfare by $M\text{-}OPT^{(R'_{\min})}_{R_{\min}}$. So, $M\text{-}OPT^{(R'_{\min})}_{R_{\min}} = S\text{-}OPT_{R'_{\min}} = \sum_{i=1}^{R'_{\min}} v'_i$, where v'_i is the valuation of *virtual bidder* i. In our spectrum auction scheme, we randomly choose a q from $\{2^1, 2^2, \ldots, 2^i, \ldots, m\}$, where m is the total demands of all bidders. It is easy to see that the probability of choosing each possible value of q is $1/\log m$. Like the optimal solution, we assume the q virtual bidders are corresponding to at most Q original bidders. Let $\mathbb{P}^{(\mathrm{E})}$ denote the set of permutations on eligible bidders, $P^{(\mathrm{E})}_k$ denote the k_{th} eligible bidder permutation and T_k denote the size of the largest set of satisfied (original) bidders in $P^{(\mathrm{E})}_k$, whose sum of demands does not exceed $R'_{\min}$. Note that, because original bidders will have different channel demands, we use $T = \min\left\{T_k|P^{(\mathrm{E})}_k \in \mathbb{P}^{(\mathrm{E})}\right\}$ to denote the minimum T_k for all possible permutations. This is also to guarantee that T original bidders could be satisfied for all possible permutations (allocations). In the following, we show it is sufficient to analyze two possible cases to derive the approximation ratio with respect to the optimal social welfare.

- Case 1: $R'_{\min} < q \leq 2R'_{\min}$. In this case, we select winners from original bidders after mapping virtual bidders to original bidders. Since the last original bidder in $R_{\min}$ bidders may be a partially-satisfied bidder, we assume one of the original bidder in Q is divided into two new original bidders. Thus, we at least select T original bidders from $Q+1$ original bidders. Due to the random selection of eligible bidders, the expectation of social welfare is greater than $\frac{T}{Q+1}M\text{-}OPT^{(q)}_Q$.
- Case 2: $\frac{1}{2}R'_{\min} < q \leq R'_{\min}$. In this case, there exist a surplus of selling items to be allocated to eligible virtual bidders. Thus, all of the eligible virtual bidders can be satisfied. The social welfare is $M\text{-}OPT^{(q)}_Q$.

We denote the social welfare of the case 1 and the case 2 by SW_1 and SW_2, respectively. In case 1, because $q > R'_{\min}$, $M\text{-}OPT^{(R'_{\min})}_{R_{\min}} = S\text{-}OPT_{R'_{\min}}$ and $M\text{-}OPT^{(q)}_Q = S\text{-}OPT_q$, we have $SW_1 \geq \frac{T}{Q+1}M\text{-}OPT^{(q)}_Q > \frac{T}{Q+1}M\text{-}OPT^{(R'_{\min})}_{R_{\min}}$. In case 2, because $q > \frac{1}{2}R'_{\min}$ and v'_i's are sorted in a non-increasing order, $SW_2 = M\text{-}OPT^{(q)}_Q > \frac{1}{2}M\text{-}OPT^{(R'_{\min})}_{R_{\min}}$. Thus, the social welfare of the proposed spectrum auction scheme *in expectation* is lower-bounded by

$$\begin{aligned} E[SW] &= (1/\log m)\cdot(\sum_{i=1}^{\log m} SW_i) \\ &\geq (1/\log m)\cdot(SW_1 + SW_2) \\ &\geq (1/\log m)\cdot(\frac{T}{Q+1}M\text{-}OPT^{(R'_{\min})}_{R_{\min}} + \frac{1}{2}M\text{-}OPT^{(R'_{\min})}_{R_{\min}}) \\ &= (1/\log m)\cdot(\frac{T}{Q+1} + \frac{1}{2})M\text{-}OPT^{(R'_{\min})}_{R_{\min}}. \end{aligned}$$

■

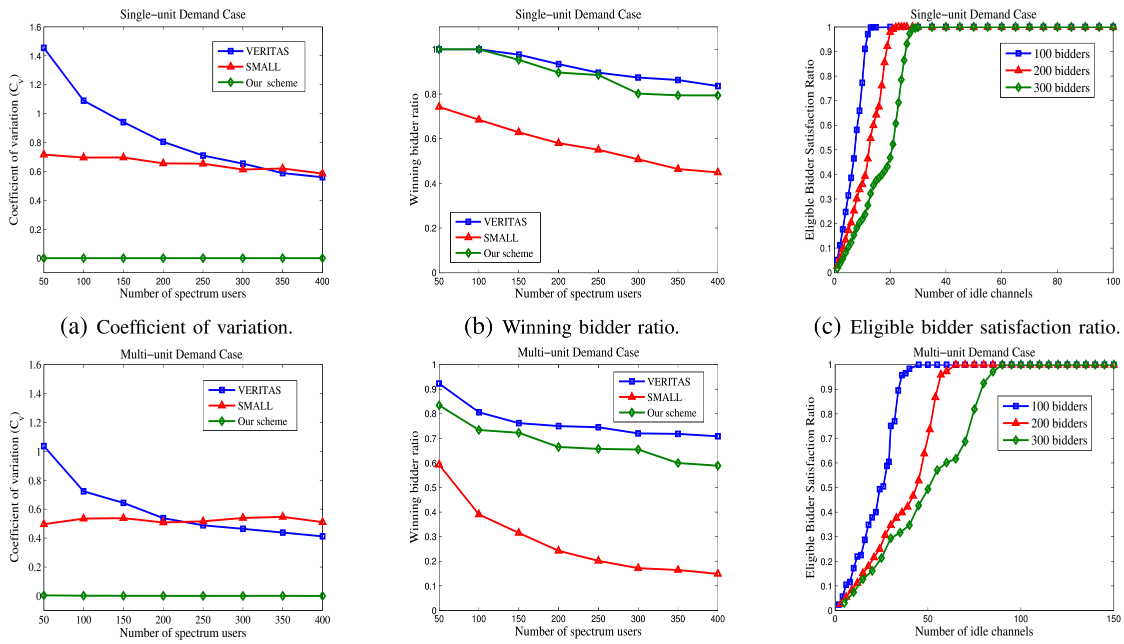

(a) Coefficient of variation. (b) Winning bidder ratio. (c) Eligible bidder satisfaction ratio.

(d) Coefficient of variation. (e) Winning bidder ratio. (f) Eligible bidder satisfaction ratio.

Fig. 5. Comparing our spectrum auction schemes to the existing benchmarks. Top 3 figures assume single-unit demand and bottom 3 are for multi-unit demand.

IV. Experimental Results

In this section, we perform experiments to evaluate the performance of our spectrum auction schemes under a dynamic supply of idle channels. We explore the unique properties of our spectrum auction schemes by comparing them to two well-known spectrum auction mechanisms VERITAS [1] and SMALL [3].

A. Methodology

In our experiments, the number of bidders varies from 50 to 400. Bidders are randomly distributed in the area of 2000×2000 square metres and the conflict range for each bidder is 400 meters. As shown in previous analysis, the efficiency performance bound derived above does not rely on the distribution of the arrival of idle channels. Thus, the use of any distribution in our experiments will not affect the results. In the single-unit case, each bidder requests one channel, while in the multi-unit case each bidder requests multiple channels. For each bidder, the per-channel bid is randomly distributed over $(0, 1]$ [1], [3]. All experimental results are averaged over 200 rounds. For bidder grouping, we divide eligible bidders into interference-free groups using the greedy grouping algorithm in [14].

The highlight of our spectrum auction mechanisms is the realization of *fair pricing* and *truthful bidding* simultaneously in a practical setting where idle channels are supplied to spectrum users in a dynamic manner. In addition, our spectrum auction designs are randomized mechanisms. To carefully characterize the unique properties, we use the following three customized performance metrics.

- *Coefficient of Variation*. Coefficient of variation is a statistical measure of the dispersion of data points in a data series around the mean, and it is defined as the ratio of the standard deviation σ to the mean μ, *i.e.*, $C_v = \frac{\sigma}{\mu}$. We compute C_v of per-channel payments of all winners to evaluate the fairness of pricing. Obviously, the lower C_v the better fairness can be achieved, and it means that there exist less differences among the per-channel payments of all winners.
- *Winning Bidder Ratio*. Winning bidder ratio is used to evaluate if all the winners are the ones that have the highest per-channel bids among all bidders, and it is defined as the ratio of the number of winners with the highest bids (among all bidders) to the number of winners. In fact, this metric reflects the fairness of allocation, which means bidders with higher per-channel bids should be satisfied when idle channels are arriving dynamically.
- *Eligible Bidder Satisfaction Ratio*. Eligible bidder satisfaction ratio is used to evaluate if all eligible bidders can be satisfied given an unknown supply of idle channels, and it is defined as the percentage of winners or say satisfied bidders among all eligible bidders.

Note that our mechanisms randomly choose the number of eligible bidders and only an eligible bidder has the opportunity to get allocated. This characteristic is distinct from the existing spectrum auction schemes such as VERITAS [1] and SMALL [3]. So, we do not compare the *eligible bidder satisfaction ratio* of our solutions to the existing ones. In addition, VERITAS and SMALL can only support single-round offline spectrum auction, so we carefully make an extension of them to obtain a fair comparison. For VERITAS, we perform channel allocation once some idle channels arrive and then do the pricing for winners. The critical value is computed for the unsatisfied bidders in previous rounds while taking into account the inference constraints caused by previous winners. For SMALL, we randomly generate channel reserve price for each coming idle channel from $[0, 2)$ following the same assumption made in [3]. When idle channels arrive, we run SMALL for all unused channels and unallocated groups.

B. Coefficient of Variation

Fig. 5 (a) and (d) compare C_vs of winners' payments in VERITAS, SMALL and our schemes against the number of spectrum users. As can be seen, despite of the increase of spectrum users our spectrum auction schemes achieve zero C_v in the single-unit demand case and *almost* zero C_v in the multi-unit demand case. This implies that in comparison to VERITAS and SMALL, our spectrum auction design indeed provide fair pricing for spectrum users bidding for homogenous channels.

C. Winning Bidder Ratio

Fig. 5 (b) and (e) compare the winning bidder ratios of VERITAS, SMALL and our schemes against the number of spectrum users. The results show that VERITAS and our schemes have similar performance, ensuring that most of bidders with higher bids will become winners. This is because VERITAS adopts a greedy allocation strategy and our schemes also choose q eligible bidders with the highest bids for future allocation. The small loss of performance is due to the randomization of eligible bidders during the channel allocation process. SMALL achieves a relatively lower ratio since it allocates idle channels based on the group bids, which relates to the bidder with the lowest bid in the group. In some cases, it is hard to get some bidders with higher bids allocated since the groups they belong to have very low group bids.

D. Eligible Bidder Satisfaction Ratio

Different from the existing spectrum auction schemes (including VERITAS and SMALL), the *winner selection* and *channel allocation* of our schemes is randomized to select so-called eligible bidders, which obtain the opportunities to get allocated. To obtain the percentage of winners or say satisfied bidders among all eligible bidders, we evaluate the eligible bidder satisfaction ratio under different idle channel supplies. In Fig. 5 (c) and (f), as expected, the larger the bidder set the larger number of idle channels is required to satisfy all eligible bidders. We can also observe that a small number of idle channels can satisfy a larger number of eligible bidders due to spectrum reusability.

V. Related Work

Spectrum auction, which allows an authority to sell licenses for signal transmission over specific bands, allocates scarce spectrum resources quickly and efficiently to the users that value them the most. In contrast to the auction of other digital goods, spectrum has a very unique characteristic called reusability due to the inherent nature of interference in radio transmissions. That is, users whose radio transmissions do not interfere each other in different geographic locations are able to share the same spectrum simultaneously. Obviously, spectrum/frequency reusability enables the communication system to increase both coverage and capacity. However, it also poses new challenges for the spectrum auction design, *e.g.*, the reusability makes it hard to achieve a truthful design and to provide bounded performance in terms of social welfare/revenue. Recently, spectrum auctions have received extensive research efforts in the literature [1]–[7]. In [1], Zhou *et al.* utilized the greedy allocation together with the critical value based pricing in [15] to design a truthful and computationally efficient spectrum auction mechanism under the bidder interference constraints. In [5], the same authors proposed Athena, a *collusion-resistant* spectrum auction mechanism using APM [16]. By first dividing bidders into multiple non-overlapping segments based on the interference constraints, the spectrum auction problem is transferred to a single-unit auction problem with an unlimited frequency/channel supply and the partition independence property of bids enables cheating-resistance. Similar models of spectrum auction [3], [4], [7] were studied and the resulting auction schemes with different design goals have been proposed in recent years. Huang et al. [4] presented a truthful and privacy-preserving spectrum auction mechanism, called SPRING, which used two cryptographic tools order-preserving symmetric encryption (OPSE) [17] and oblivious transfer (OT) [18] to prevent attackers from learning private information of bidders. Wu et al. [3] proposed a mechanism, called SMALL, to adapt to bidders with multiple radios by dividing bidders into segments and setting a reserve price for channels such that bidders cannot benefit by manipulating their own bids. These models addressed the problem of spectrum reusability by using the graph coloring algorithm to divide bidders into different groups [8], and the smallest bid is selected in each group as a *group bid* for resisting bidder misreport. In [7], Chen et al. proposed TAMES, a truthful double auction for bidders with heterogeneous demands and showed that the auctioneer can choose different *graph coloring algorithms* for different goals. While theoretically sound, these spectrum auction mechanisms do not provide performance bounds to characterize the system performance in terms of social welfare or revenue. In addition, they only focus on the offline auction model, where the set of users (*i.e.*, channel bidders) and the set of goods (*i.e.*, channels) are pre-determined before the start of auction process.

Recently, researches on online spectrum auction models have aroused much interest [9]–[11]. In general, there are two types of *online* spectrum auction models: spectrum users arrive in an online manner and spectrum supplies arrive in an online manner. In [9], Deek et al. made an extension of [1] and investigated the online multi-good selling scenario in [19]. Considering the online arrival of users, the authors proposed an efficient online spectrum auction mechanism with *preemption*, enabling the resistance of both bid- and time-based cheating. However, it does not provide a performance bound on revenue with respect to the optimal solution in general. Under the same online auction model, Xu et al. [10] proposed TOFU, another online semi-truthful spectrum auction scheme with channel *preemption*. To carefully characterize the auction model, the authors derived *competitive ratios* under different application scenarios. Different from the previous solutions, TOFU achieves only semi-truthfulness, where users are able to underbid to gain self-benefits. In practice, while the set of bidders is known to the auctioneer, the exact number of items for sale may be uncertain, *e.g.*, the spectrum auctioneer allows users to bid for idle channels released dynamically by other spectrum users. Thus, another *online* model involves the online arriving spectrum supplies during auction. However, online spectrum allocation design on this model has received limited attention so far. Besides, the dynamic release of channels to users also arouses the question of designing a *fair* pricing scheme. Almost all existing spectrum auction models consider homogeneous channels, but the payment for a channel will be quite different. From the perspective of users, it is unfair for them to make different payments for the same goods especially when idle channels are arriving online. To solve this problem, a line of studies on *envy-free* auction design has been proposed [20]–[23], where all users can obtain the same goods by the same payment. These studies focus on different goals, including maximizing the total profit [20], proving the competitive ratio [21], pricing optimally on single-minded combinatorial auctions [22] and considering distributional information in the auction design [23].

VI. Concluding Remarks and Future Work

In this paper, we investigated the problem of allocating channels to spectrum users in a setting where the supply of idle channels are unknown, with the goal of maximizing social welfare. We designed a suite of truthful and efficient spectrum auction schemes, providing different users with fair pricing for homogenous channels and guarantee bounded performance with respect to the optimal social welfare. We analytically showed the truthfulness and approximation ratio of our spectrum auction designs. Our experimental results validated our analysis and demonstrated the desirable properties of our proposed designs.

To our best knowledge, our current work represents the first effort to formulate and investigate the truthful frequency allocation problem with a dynamic spectrum supply. Several important open problems remain for future research. First, in our spectrum auction scheme for the multi-unit demand case, to achieve truthful auction design and maximize the social welfare, we weaken the strict demand requirement by assuming that spectrum users are not strictly single-minded. In our future work, it is important to explore the case where all users are single-minded, *i.e.*, a user is willing to pay only if it is fully-satisfied. We still aim to achieve both efficiency and truthfulness for the proposed mechanisms. Second, in many case studies of auctions run in practice, collusion is considered as a serious problem. Thus, we propose to design collusion-resistant mechanisms for online spectrum auction with a dynamic supply while achieving fair pricing, strategy-proofness and bounded performance.

Acknowledgment

Kui's research is supported in part by US National Science Foundation under grants CNS-1262275 and CNS-1318948. Qian's research is supported in part by National Natural Science Foundation of China (Grant No. 61373167) and Natural Science Foundation of Hubei Province (Grant No. 2013CFB297).

References

[1] X. Zhou, S. Gandhi, S. Suri, and H. Zheng, "ebay in the sky: strategy-proof wireless spectrum auctions," in *Proceedings of the 14th ACM international conference on Mobile computing and networking*. ACM, 2008, pp. 2–13.

[2] Y. Chen, J. Zhang, K. Wu, and Q. Zhang, "Tames: A truthful auction mechanism for heterogeneous spectrum allocation," *INFOCOM, 2013 Proceedings IEEE*, pp. 180–184, 2013.

[3] F. Wu and N. Vaidya, "A strategy-proof radio spectrum auction mechanism in noncooperative wireless networks," *IEEE Transactions on Mobile Computing*, vol. 12, no. 5, pp. 885–894, 2013.

[4] Q. Huang, Y. Tao, and F. Wu, "Spring: A strategy-proof and privacy preserving spectrum auction mechanism," *INFOCOM, 2013 Proceedings IEEE*, pp. 827–835, 2013.

[5] X. Zhou and H. Zheng, "Breaking bidder collusion in large-scale spectrum auctions," in *Proceedings of the eleventh ACM international symposium on Mobile ad hoc networking and computing*. ACM, 2010, pp. 121–130.

[6] J. Jia, Q. Zhang, Q. Zhang, and M. Liu, "Revenue generation for truthful spectrum auction in dynamic spectrum access," in *Proceedings of the tenth ACM international symposium on Mobile ad hoc networking and computing*. ACM, 2009, pp. 3–12.

[7] Y. Chen, J. Zhang, K. Wu, and Q. Zhang, "Tames: A truthful double auction for multi-demand heterogeneous spectrums," *IEEE Transactions on Parallel and Distributed Systems*, vol. 99, no. PrePrints, p. 1, 2013.

[8] B. Douglas, "West. introduction to graph theory," 1996.

[9] L. Deek, X. Zhou, K. Almeroth, and H. Zheng, "To preempt or not: Tackling bid and time-based cheating in online spectrum auctions," *INFOCOM, 2011 Proceedings IEEE*, pp. 2219–2227, 2011.

[10] P. Xu and X.-Y. Li, "Tofu: Semi-truthful online frequency allocation mechanism for wireless networks," *Networking, IEEE/ACM Transactions on*, vol. 19, no. 2, pp. 433–446, 2011.

[11] P. Xu, X. Xu, S. Tang, and X.-Y. Li, "Truthful online spectrum allocation and auction in multi-channel wireless networks," *INFOCOM, 2011 Proceedings IEEE*, pp. 26–30, 2011.

[12] D. Pisinger, "Algorithms for knapsack problems," 1995.

[13] W. Vickrey, "Counterspeculation, auctions, and competitive sealed tenders," *The Journal of finance*, vol. 16, no. 1, pp. 8–37, 1961.

[14] D. J. Welsh and M. B. Powell, "An upper bound for the chromatic number of a graph and its application to timetabling problems," *The Computer Journal*, vol. 10, no. 1, pp. 85–86, 1967.

[15] D. Lehmann, L. I. Oćallaghan, and Y. Shoham, "Truth revelation in approximately efficient combinatorial auctions," *Journal of the ACM (JACM)*, vol. 49, no. 5, pp. 577–602, 2002.

[16] A. V. Goldberg and J. D. Hartline, "Collusion-resistant mechanisms for single-parameter agents," in *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 2005, pp. 620–629.

[17] R. Agrawal, J. Kiernan, R. Srikant, and Y. Xu, "Order preserving encryption for numeric data," in *Proceedings of the 2004 ACM SIGMOD international conference on Management of data*. ACM, 2004, pp. 563–574.

[18] W.-G. Tzeng, "Efficient 1-out-of-n oblivious transfer schemes with universally usable parameters," *Computers, IEEE Transactions on*, vol. 53, no. 2, pp. 232–240, 2004.

[19] M. T. Hajiaghayi, "Online auctions with re-usable goods," in *Proceedings of the 6th ACM conference on Electronic commerce*. ACM, 2005, pp. 165–174.

[20] V. Guruswami, J. D. Hartline, A. R. Karlin, D. Kempe, C. Kenyon, and F. McSherry, "On profit-maximizing envy-free pricing," in *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 2005, pp. 1164–1173.

[21] A. V. Goldberg and J. D. Hartline, "Envy-free auctions for digital goods," in *Proceedings of the 4th ACM conference on Electronic commerce*. ACM, 2003, pp. 29–35.

[22] J. Hartline and Q. Yan, "Envy, truth, and profit," in *Proceedings of the 12th ACM conference on Electronic commerce*. ACM, 2011, pp. 243–252.

[23] M. Babaioff, L. Blumrosen, and A. Roth, "Auctions with online supply," in *Proceedings of the 11th ACM conference on Electronic commerce*. ACM, 2010, pp. 13–22.