Truthful Spectrum Auction Design for Secondary Networks

Yuefei Zhu, Baochun Li
Department of Electrical and Computer Engineering
University of Toronto

Zongpeng Li
Department of Computer Science
University of Calgary

***Abstract*—Opportunistic wireless channel access by non-licensed users has emerged as a promising solution for addressing the bandwidth scarcity challenge. Auctions represent a natural mechanism for allocating the spectrum, generating an economic incentive for the licensed user to relinquish channels. A severe limitation of existing spectrum auction designs lies in the over-simplifying assumption that every non-licensed user is a single-node or single-link *secondary user*. While such an assumption makes the auction design easier, it does not capture practical scenarios where users have multihop routing demands. For the first time in the literature, we propose to model non-licensed users as *secondary networks* (SNs), each of which comprises of a multihop network with end-to-end routing demands. We aim to design truthful auctions for allocating channels to SNs in a coordinated fashion that maximizes social welfare of the system. We use simple examples to show that such auctions among SNs differ drastically from simple auctions among single-hop users, and previous solutions suffer severely from local, per-hop decision making. We first design a simple, heuristic auction that takes inter-SN interference into consideration, and is truthful. We then design a randomized auction based on primal-dual linear optimization, with a proven performance guarantee for approaching optimal social welfare. A key technique in our solution is to decompose a linear program (LP) solution for channel assignment into a set of integer program (IP) solutions, then applying a pair of tailored primal and dual LPs for computing probabilities of choosing each IP solution. We prove the truthfulness and performance bound of our solution, and verify its effectiveness through simulation studies.**

I. Introduction

Recent years have witnessed substantial growth in wireless technology and applications, which rely crucially on the availability of bandwidth spectrum. Traditional spectrum allocation is static, and is prone to inefficient spectrum utilization in both temporal and spatial domains: large spectrum chunks remain idling while new users are unable to access them. Such an observation has prompted research interest in designing a secondary spectrum market, where new users can access a licensed channel when not in use by its owner, with appropriate remuneration transferred to the latter.

In a secondary spectrum market, a spectrum owner or *primary user* (PU) leases its idle spectrum chunks (channels) to *secondary users* (SUs) through auctions [1], [2]. SUs submit bids for channels, and pay the PU a price to access a channel if their bids are successful. A natural goal of spectrum auction design is *truthfulness*, under which an SU's best strategy is to bid its true valuation of a channel, with no incentive to lie. A truthful auction simplifies decision making at SUs, and lays a foundation for good decision making at the PU. Another important goal in spectrum auction design is *social welfare maximization*, *i.e.*, maximizing the aggregated 'happiness' of everyone in the system. Such an auction tends to allocate channels to SUs who value them the most.

A unique feature of spectrum auction design is the need of appropriate consideration for wireless interference and spatial reuse of channels. A channel can be allocated to multiple SUs provided that they are far apart, with no mutual interference. Optimal channel assignment for social welfare maximization is equivalent to the graph colouring problem, and is NP-hard [3], even assuming truthful bids are given for free. Existing works on spectrum auctions mostly focus on resolving such a challenge (*e.g.*, [4], [5]) while assuming the simplest model of a SU: a single node, or a single link, similar to a single hop transmission in cellular networks [2], [4], [5].

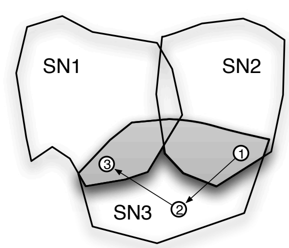

Fig. 1: A secondary spectrum market with 3 SNs and 2 channels.

After extensive research during the past five years, auction design for single-hop users, each requesting a single channel, has been relatively well understood. However, a practical SU may very well comprise of multiple nodes forming a multi-hop network, which we refer to as a *secondary network* (SN). These include scenarios such as users with multihop access to base stations, or users with their own mobile ad hoc networks. SNs require coordinated end-to-end channel assignment, and in general benefit from multi-channel diversity along its path. The SN model subsumes the SU model as the simplest special case.

Fig. 1 depicts three co-located SNs, SN1, SN2 and SN3, which have interference with one another, because their network regions overlap. The primary network (PN) has two channels, Ch1 and Ch2, which have been allocated to SN1 and SN2, respectively. Now SN3 wishes to route along a two-hop path $1 \rightarrow 2 \rightarrow 3$. Under existing single-channel auctions for SUs, SN3 cannot obtain a channel, because each channel interferes with either SN1 or SN2. However, a solution exists by relaxing the one channel per user assumption, and assigning Ch1 to link $1 \rightarrow 2$ and Ch2 to the link $2 \rightarrow 3$. In general, taking multichannel, multihop transmissions by SNs into consideration can apparently improve channel utilization and social welfare. Note here that the model in which an SN bids for multiple channels is inapplicable, because due to the unawareness of other SNs' information, an SN cannot know the number of channels to bid for, to form a feasible path.

Designing truthful auctions for SNs is an interesting problem, but by no means an easy one. We note that it is hard for an SN to decide by itself an optimal or good path to bid for. Such decision making requires global information on other SNs as well, and is naturally best made by the auctioneer, *i.e.*, the PN. Consequently, a bid from an SN includes just a price it wishes to pay, with two nodes it wishes to connect using a path. Furthermore, SNs now interfere with each other in a more complex manner. Not only that they transmit along multihop paths, but each path can be assigned with distinct channels at different links. The PN, after receiving bids, needs to make judicious joint routing and channel assignment decisions.

In this work, we first design a simple heuristic auction for spectrum allocation to SNs, which guarantees both truthfulness and interference-free channel allocation, providing winning SNs with end-to-end multihop paths, with a channel assigned to each hop. The heuristic auction enables multi-channel assignment along a path, thereby reducing the possibility that a path is blocked due to interference. To achieve truthfulness, we employ a greedy, monotonic allocation rule and design an accompanying payment scheme, by referring to Myerson's characterization of truthful auctions [6].

The heuristic auction provides no hard guarantee on social welfare. We next design a randomized auction, which is truthful in expectation, and is provably approximate optimal in social welfare. We note that absolute optimal social welfare is impossible, since the joint routing-channel assignment problem is already NP-hard with truthful bids given for free. We relax an integer program (IP) formulation to the social welfare maximization problem into a linear program (LP), and prove an upper-bound on the integrality gap. We then employ the decomposition technique due to Lavi and Swamy [7] to decompose an LP solution into a set of feasible IP solutions (allocations). A pair of primal-dual LPs are formulated, for computing a probability distribution over the allocation set. Based on the set and the probabilities, an approximation algorithm is finally designed, for computing a feasible solution to the original IP with a provable performance guarantee. Such a solution assigns channels to paths constructed for winning SNs. We then apply the classic VCG [8]–[10] payment scheme, to conclude the design of a truthful auction that approximately maximizes social welfare for spectrum allocation to SNs.

The remainder of the paper is organized as follows. We discuss related work in Sec. II, and present preliminaries in Sec. III. A heuristic truthful auction is designed in Sec. IV. In Sec. V, we propose and analyze a randomized auction with performance bound. Simulation studies are presented in Sec. VI. Sec. VII concludes the paper.

II. Related Work

Auctions serve as an efficient mechanism for distributing scarce resources to competing participants in a market. To simplify the strategical behaviour of agents and hence encourage participation, truthfulness is desired. A celebrated work is the VCG mechanism due to Vickrey [10], Clarke [8], and Groves [9]. However, the VCG mechanism is only suitable when optimal solutions are computationally feasible, and is not directly applicable for secondary spectrum auctions, because interference-free channel allocation is NP-Hard.

Allocating spectrum using auctions has received considerable research attention recently. Early solutions include auctions that allocate power [11] and allocate a channel to each winning user [12]. These auctions are unfortunately not truthful. VERITAS [1] is the first truthful spectrum auction based on the monotonic allocation rule. Zhou and Zheng propose TRUST [13], which is a double auction with multiple sellers (licensed users). Jia *et al.* [14] design a spectrum auction mechanism that not only encourages truthful behaviour but also computes approximately maximum revenue, which is an alternative goal to maximum social welfare.

For spectrum auctions that take interference among secondary users into consideration, Wu *et al.* [2] develop a semi-definite programming based mechanism, which is truthful and resistant to bidder collusion. Gopinathan [5] *et al.* propose auctions that incorporate fairness considerations into channel allocation. Their goal is to maximize social welfare, while ensuring a notion of fairness among bidders when the auction is repeatedly held. A truthful and scalable spectrum auction enabling both sharing and exclusive access is proposed by Kash *et al.* [4]. This auction handles heterogeneous agent types with different transmission powers and spectrum needs. We note that, all the works mentioned above focus on single-hop users bidding for a single channel only. Our work essentially generalizes the problem to multi-hop users, who may enjoy multi-channel paths.

III. Preliminaries

In this section, we first introduce some background in truthful auction design in Sec. III-A, then describe our system model in Sec. III-B.

A. Truthful Auction Design

Auction theory is a branch of economics that studies how people act in an auction and analyzes the properties of auction markets. We first introduce some basic and most related concepts, definitions and theorems from auction design.

An auction allocates items or goods (channels in our case) to competitive agents with bids and private valuations. We adopt w_i as nonnegative valuations of each agent i, which is often private information known only to the agent itself. Besides determining an allocation, an auction also computes payments/charges for winning bidders. We denote by $p(i)$ and b_i the payment and bid of agent i, respectively. Then the utility of i is a function of all the bids:

$$u_i(b_i, \mathbf{b}_{-i}) = \begin{cases} w_i - p(i) & \text{if agent } i \text{ with bid } b_i \text{ gets an item} \\ 0 & \text{otherwise} \end{cases}$$

where $\mathbf{b}_{-i}$ is the vector of all the bids except b_i. We first adopt some conventional assumptions in economics here. We assume that each agent i is selfish and rational. A *selfish* agent is one that acts strategically to maximize its utility. An agent is said to be *rational* in that it always prefers the outcome that brings itself a larger utility. Hence, an agent i may lie about its valuation, and bid $b_i \neq w_i$ if doing so yields a higher utility.

Truthfulness is a desirable property of an auction, where reporting true valuation in the bid is optimal for each agent i, regardless of other agents' bids. If agents have incentives to lie, other agents are forced to strategically respond to these lies, making the auction and its analysis complex. A key advantage of a truthful auction is that it simplifies agent strategies. Formally, an auction is *truthful* if for any agent i with any $b_i \neq w_i$, any $\mathbf{b}_{-i}$, we have

$$u_i(w_i, \mathbf{b}_{-i}) \geq u_i(b_i, \mathbf{b}_{-i}) \quad (1)$$

An auction is *randomized* if its allocation decision making involves flipping a (biased) coin. The payment and utility of an agent are then random variables. A randomized auction is *truthful in expectation* if (1) holds in expectation. Besides, we also prefer an auction to be *individually rational*, in which agents pay no more than their gain (valuations).

As discussed, the classic VCG mechanism for truthful auction design requires the optimal allocation to be efficiently computable, and is not practical for spectrum auctions, since optimal channel allocation is NP-hard. If we aim to design a tailored, heuristic truthful auction, then we may rely on the characterization of truthful auctions by Myerson [6].

Theorem 1. *Let $P_i(b_i)$ be the probability of agent i with bid b_i winning an auction. An auction is truthful if and only if the followings hold for a fixed $\mathbf{b}_{-i}$:*

- *$P_i(b_i)$ is monotonically non-decreasing in b_i;*
- *Agent i bidding b_i is charged $b_i P_i(b_i) - \int_0^{b_i} P_i(b)\mathrm{d}b$.*

Given Theorem 1, we see that once the allocation rule $\mathbf{P}(\cdot) = \{P_i(b_i)\}_{i\in\mathcal{N}}$ is fixed ($\mathcal{N}$ is the set of bidders), the payment rule is also fixed. For the case where the auction is deterministic, there are two equivalent ways to interpret Theorem 1: (i) there exists a minimum bid b_i^*, such that i will win only if agent i bids at least b_i^*, *i.e.*, the monotonicity of $P_i(b_i)$ implies that, there is some *critical bid* b_i^*, such that $P_i(b_i)$ is 1 for all $b_i > b_i^*$ and 0 for all $b_i < b_i^*$; (ii) the payment charged to agent i for a fixed $\mathbf{b}_{-i}$ should be independent of b_i (formally, $p_i(b_i) = b_i - \int_{b_i^*}^{b_i} \mathrm{d}b = b_i^*$).

B. System Model

We assume there is a set of SNs, $\mathcal{N}$. Each SN has deployed a set of nodes in a geographical region, and has a demand for multihop transmission from a source to a destination. A PN has a set of channels, $\mathcal{C}$, available for auctioning in the region. We refer to SNs as *agents* and the PN as the *auctioneer*. Each node within an SN is equipped with a radio that is capable of switching between different channels. SNs do not collaborate with each other, and nodes from different SNs are not required to forward traffic for each other.

TABLE I: List of notations

w_i	valuation of agent i	b_i	bid of agent i
$\mathbf{b}_{-i}$	bid of all agents except b_i	$\phi(i)$	virtual bid of agent i
$p(i)$	payment of agent i	u^i	a node in SN i
l_{uv}^i	link from u^i to v^i	f_{uv}^i	flow rate on link l_{uv}^i
$O(\mathbf{w})$	objective function of IP (5)	$S(\mathbf{w})$	objective function of the LPR
$I_s(i)$	the set of SNs that interfere with SN i along its path		
$x(c, l_{uv}^i)$	binary var: whether channel c is allocated to link l_{uv}^i		
$G^i(\mathcal{E}^i, \mathcal{V}^i)$	connectivity graph of SN i with link set $\mathcal{E}^i$, node set $\mathcal{V}^i$		
$H(\mathcal{E}_H, \mathcal{V}_H)$	conflict graph of links of all the SNs with edge set $\mathcal{E}_H$ and vertex set $\mathcal{V}_H$		

We assume nodes from each SN i form a connected graph $G^i(\mathcal{E}^i, \mathcal{V}^i)$, which also contains node locations. We use "node" and "link" for the connectivity graphs and "vertex" and "edge" for the conflict graph introduced later. To better formulate the joint routing-channel assignment problem, we incorporate the concept of network flows. Let u^i be a node in SN i and s^i, d^i be the source and the destination in SN i. We use l_{uv}^i to denote the link from node u^i to node v^i belonging to SN i, and f_{uv}^i to denote the amount of flow on link l_{uv}^i. Later we connect d^i back to s^i with a virtual feedback link l_{ds}^i, for a compact formulation of the joint optimization IP.

We define a *conflict graph* $H(\mathcal{E}_H, \mathcal{V}_H)$, whose vertices correspond to links from all the connectivity graphs. We use (l_{uv}^i, l_{pq}^j) to denote an edge in $\mathcal{E}_H$, indicating that link l_{uv}^i and link l_{pq}^j interfere if allocated a common channel. Before the auction starts, each SN i submits to the auctioneer a compound bid, defined as $\mathfrak{B}_i = (G^i(\mathcal{E}^i, \mathcal{V}^i), s^i, d^i, b_i)$. Then the conflict graph can be centrally obtained by the auctioneer. We denote by w_i the private valuation of SN i for a feasible path between s^i and d^i, and $p(i)$ its payment. b_i, w_i and $p(i)$ all represent monetary amounts. Note that we assume agents only have incentives to lie about their valuations.

We denote by $R_T(u^i)$ and $R_I(u^i)$ the transmission range and interference range of node u^i, respectively. We assume that $\frac{R_I(u^i)}{R_T(u^i)} = \Delta$ and $R_T(u^i) \leq R_{\max}$ for any node u^i where $\Delta \geq 1$. Since no inter-SN collaboration is assumed, links from different SNs do not participate in joint MAC scheduling, and cannot be assigned the same channel if they interfere. As a result, two links l_{uv}^i and l_{pq}^j interfere if a node in $\{u, v\}$ is within the interference range of a node in $\{p, q\}$, and cannot be assigned the same channel if $i \neq j$. Formally, let a binary variable $x(c, l_{uv}^i) \in \{0, 1\}$ denote whether channel $c \in \mathcal{C}$ is assigned to link l_{uv}^i for user i. If for channel $c \in \mathcal{C}$, $x(c, l_{uv}^i) = x(c, l_{pq}^j)$, then $(l_{uv}^i, l_{pq}^j) \notin \mathcal{E}_H$. Hence, for the joint routing-channel assignment problem we have the **Channel Interference Constraints**:

$$x(c, l_{uv}^i) + x(c, l_{pq}^j) \leq 1, \;\; (l_{uv}^i, l_{pq}^j) \in \mathcal{E}_H, \;\; \forall c \in \mathcal{C} \quad (2)$$

We also need **Flow Conservation Constraints**, *i.e.*, at any node in $\mathcal{V}^i$, the total incoming and outgoing flows equal (recall the virtual feedback link):

$$\sum_{u \in \mathcal{V}^i} f^i_{uv} = \sum_{u \in \mathcal{V}^i} f^i_{vu}, \;\; \forall v \in \mathcal{V}^i \tag{3}$$

Assuming each channel has the same unit capacity 1, we have the **Capacity Constraints**:

$$\sum_{u \in \mathcal{V}^i \setminus \{d^i\}} f^i_{uv} \leq \sum_{c \in \mathcal{C}} x(c, l^i_{uv}) \leq 1 \tag{4}$$

which also ensures that a link can be assigned a single channel only.

An agent needs an end-to-end path between its source and destination. This corresponds to a network flow of rate 1. Note that the link flow on the feedback link f^i_{ds} equals the end-to-end flow for SN i. We formulate the joint routing-channel assignment problem for SNs into an IP:

$$\text{maximize } O(\mathbf{w}) = \sum_{i \in \mathcal{N}} w_i f^i_{ds} \tag{5}$$

subject to

$$\begin{aligned}
& x(c, l^i_{uv}) + x(c, l^j_{pq}) \leq 1, && (l^i_{uv}, l^j_{pq}) \in \mathcal{E}_H, \forall c \in \mathcal{C} \\
& \sum_{u \in \mathcal{V}^i} f^i_{uv} = \sum_{u \in \mathcal{V}^i} f^i_{vu}, && \forall v \in \mathcal{V}^i \\
& \sum_{u \in \mathcal{V}^i \setminus \{d^i\}} f^i_{uv} \leq \sum_{c \in \mathcal{C}} x(c, l^i_{uv}) \leq 1, && \forall v \in \mathcal{V}^i \\
& f^i_{uv}, \;\; x(c, l^i_{uv}) \in \{0, 1\}.
\end{aligned}$$

where $O(\mathbf{w})$ denotes the objective function of the IP. Solving this IP to optimal is an NP-hard problem. Therefore we first introduce a heuristic auction in Sec. IV, which is based on the technique of monotonic allocation and critical bids, and is simple and truthful. However, it does not provide any bound on the social welfare generated. A more sophisticated, randomized auction with a proven bound is studied next, where the LP relaxation of IP (5) is solved as a first step.

IV. A Heuristic Truthful Auction

In this section, we design an auction with a greedy style allocation and a payment scheme to ensure truthfulness. The auction consists of two phases: Algorithm 1 determines the channel assignment and winning bidders, and Algorithm 2 computes the payments for winning agents.

A. Channel Allocation

As discussed in Sec. III, the key to designing a truthful auction is to have a non-decreasing allocation rule. Prices can then be calculated by the critical bids to make the auction truthful. A greedy allocation is adopted in Algorithm 1. Assume channels are indexed by $1, 2, ...|\mathcal{C}|$. For a simple heuristic auction, we first compute the shortest path for each agent as its end-to-end path. Let $I_s(i)$ be the set of SNs that interfere with i along the path. We define the virtual bid of SN i as

$$\phi(i) = \frac{b_i}{|I_s(i)|} \tag{6}$$

The rationale behind scaling the bid by $|I_s(i)|$ is to take i's interference with other agents into consideration, for heuristically maximizing social welfare. Then we greedily assign minimum indexed available channels along the paths to each link, according to a non-increasing order of virtual bids $\phi(i)$.

Algorithm 1 *A greedy truthful auction — channel allocation.*

Input: Set of channels $\mathcal{C}$, all the compound bids $\mathfrak{B}_i = (G^i(\mathcal{E}^i, \mathcal{V}^i), s^i, d^i, b_i)$, conflict graph $H(\mathcal{E}_H, \mathcal{V}_H)$
for all $i \in \mathcal{N}$ **do**
  $I_s(i) \Leftarrow \emptyset$;
  Compute the shortest path P^i from s^i to d^i;
for all $i \in \mathcal{N}$ **do**
  for all l^i_{uv} along path P^i **do**
    $x(c, l^i_{uv}) \Leftarrow 0 \; \forall c \in \mathcal{C}$;
    if $(l^i_{uv}, l^j_{pq}) \in \mathcal{E}_H$ **then**
      $I_s(i) \Leftarrow I_s(i) \cup \{i\}$;
  $\phi(i) \Leftarrow \frac{b_i}{|I_s(i)|}$;
  $\mathrm{Win}(i) \Leftarrow \mathrm{TRUE}$;
for $i \in \mathcal{N}$ in non-increasing order of $\phi(i)$ **do**
  for all l^i_{uv} along path P^i **do**
    Let $\mathcal{T}^i_{uv} \Leftarrow \mathcal{C}$;
    for all $c \in \mathcal{T}^i_{uv}$ **do**
      if $x(c, l^j_{pq}) = 1$ with $(l^i_{uv}, l^j_{pq}) \in \mathcal{E}_H, \forall p,\; q$ **then**
        $\mathcal{T}^i_{uv} \Leftarrow \mathcal{T}^i_{uv} \setminus \{c\}$;
    if $\mathcal{T}^i_{uv} = \emptyset$ **then**
      $\mathrm{Win}(i) \Leftarrow \mathrm{FALSE}$;
  if $\mathrm{Win}(i)$ = TRUE **then**
    for all l^i_{uv} along path P^i **do**
      Choose the minimum indexed channel c_m in $\mathcal{T}^i_{uv}$;
      $x(c_m, l^i_{uv}) \Leftarrow 1$;

Fig. 2 shows an example to illustrate the channel assignment procedure. There are four SNs, a, b, c and d, where $\phi(a) > \phi(b) > \phi(c) > \phi(d)$. Two channels are available for allocation. In the figure, two intersecting links also interfere with each other. If two links from two different SNs intersect, they cannot be allocated with the same channel. The algorithm first assigns Channel 1 to SN a. As a result, it cannot assign Channel 1 to the first link of SN b, which receives Channel 2 instead, as shown in Fig. 2b, leaving SN c without a channel — it is impossible to assign either channel to c's first link. However, SN d wins, and receives a channel assignment along its path without introducing interference to a or b.

We now prove that the greedy auction is monotone.

Lemma 1. *Algorithm 1 is monotone. That is, the probability of bidder i with bid b_i winning the auction is non-decreasing in b_i, and critical bids for winning agents exist.*

Proof: Bidding higher can only increase an agent's virtual bid, and therefore increase its rank in Algorithm 1. Hence, the probability of assigning a channel to the agent is non-decreasing. Besides, Algorithm 1 is deterministic, so a critical bid b^*_i exists for a winning bidder i, such that agent i always wins if it bids $b_i \geq b^*_i$. ■

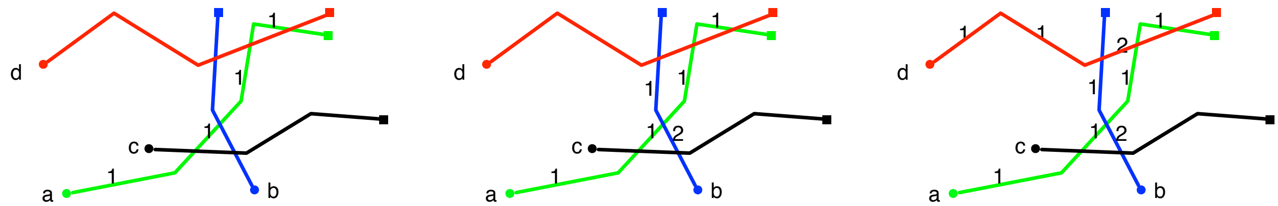

(a) Assign channels to SN a (b) Assign channels to SN b (c) Assign channels to SN d

Fig. 2: Procedure of channel assignment. Dots and squares represent source and destination nodes respectively.

B. Payment Calculation

Algorithm 2 computes payments for the winning agents. The payment scheme design is where we ensure the truthfulness of an auction. Algorithm 2 aims to find a critical bidder with critical bid b_i^* for a winning agent, such that i is guaranteed to win as long as i's virtual bid $\phi(i) \geq \phi^*(i)$. Here $\phi^*(i) = \frac{b_i^*}{|I_s(i)|}$ is the critical virtual bid for i. If b_i^* is independent from b_i, then charging agent i b_i^* will ensure that the auction is truthful, which we will argue formally later.

We now explain how Algorithm 2 works. It first clears a winning agent i's bid, and hence its virtual bid, to 0. Then we run Algorithm 1 based on $(0, \mathbf{b}_{-i})$. In Algorithm 1, an agent loses only if a link along its shortest path is unable to receive any channel. If we are unable to accommodate agent i, there must exist at least one link along its shortest path whose neighbouring links (neighbouring vertices in the conflict graph) have used all the channels. From all the agents that block links of agent i, we find out an agent j with the minimum virtual bid, set it as i's critical bidder, and compute i's payment. We claim that $\phi(i) \geq \phi(j)$, because otherwise agent i would not be a winning agent among agents in $I_s(i) \cup \{i\}$. Agent i's payment can be computed as follows:

$$p(i) = \phi^*(i)|I_s(i)| = \phi(j)|I_s(i)| \tag{7}$$

For the example in Fig. 2, we first set SN a's bid to 0, and run Algorithm 1 based on the new bid vector. After assigning channels to agent c, we find that there are no available channels for the second link of agent a. Hence, agent c becomes the critical bidder of agent a, which leads to a's payment $p(a) = \phi(c)|I_s(a)|$. The rule applies to the other two winning agents b and d as well, where $p(b) = \phi(c)|I_s(b)|$ and $p(d) = 0$.

We next show that the auction is individually rational and truthful.

Lemma 2. *The auction shown in Algorithms 1 and 2 is individually rational.*

Proof: Assume agent i wins by bidding b_i, and let j be the critical bidder of i. Then we have $\phi(i) \geq \phi(j)$, so $p(i) = \phi(j)|I_s(i)| \leq \phi(i)|I_s(i)| = b_i$. ■

Theorem 2. *The auction in Algorithms 1 and 2 is truthful.*

Proof: Fix i, $\mathbf{b}_{-i}$. Let w_i and w_i' be agent i's bid when being truthful and not, respectively. We need to show that for agent i with valuation w_i, the utility of bidding w_i is no less than the utility of bidding w_i'. We analyze the auction case by case. Let

$$r(b_i) = \begin{cases} 1 & \text{if agent } i \text{ with bid } b_i \text{ receives a channel;} \\ 0 & \text{if agent } i \text{ with bid } b_i \text{ doesn't receive a channel.} \end{cases}$$

First, assume that $w_i' < w_i$. According to Lemma 1, it is impossible for agent i to have $r(w_i) = 0$ and $r(w_i') = 1$. If $r(w_i) = 0$ and $r(w_i') = 0$, there is no incentive to lie. If $r(w_i) = 1$ and $r(w_i') = 0$, because the auction is individually rational, i also has no incentive to lie. If $r(w_i) = 1$ and $r(w_i') = 1$, then by observing that the critical bidder does not change, the utility for agent i remains the same.

Next, assume that $w_i' > w_i$. Again it is impossible for agent i to have $r(w_i) = 1$ and $r(w_i') = 0$. If $r(w_i) = 1$ and $r(w_i') = 1$, the critical bidder does not change, so the utility remains the same. If $r(w_i) = 0$ and $r(w_i') = 0$, then i has no incentive to lie. If $r(w_i) = 0$ and $r(w_i') = 1$, then at least one link is unable to receive a channel (blocked), when agent i bids w_i. From all the agents that block links of i, we find out the winning agent j with minimum virtual bid. If i bids $w_i' > w_i$ and get assigned a channel, agent j would be the critical agent of i. Then agent i's payment would be $p(i) = \phi(j)|I_s(i)| \geq w_i$, which implies a nonpositive utility of i. ■

This heuristic auction, aiming at achieving a monotonic

Algorithm 2 *A greedy truthful auction — payment calculation.*

Input: Set of channels $\mathcal{C}$, all the compound bids $\mathfrak{B}_i = (G^i(\mathcal{E}^i, \mathcal{V}^i), s^i, d^i, b_i)$, conflict graph $H(\mathcal{E}_H, \mathcal{V}_H)$, all the routing paths P^i and channel assignment from Algorithm 1.
for $i \in \mathcal{N}$ in non-increasing order of $\phi(i)$ **do**
  $p(i) \Leftarrow 0$;
  if $\mathrm{Win}(i) = 1$ **then**
    Set $b_i' \Leftarrow 0$;
    Run Algorithm 1 on $(b_i', \mathbf{b}_{-i})$;
    if Win(i) = FALSE **then**
      Let $\phi^*(i) \Leftarrow +\infty$;
      for all l_{uv}^i along path P^i **do**
        Let $\mathcal{T}_{uv}^i \Leftarrow \mathcal{C}$;
        for all $c \in \mathcal{T}_{uv}^i$ **do**
          if $x(c, l_{pq}^j) = 1$ with $(l_{uv}^i, l_{pq}^j) \in \mathcal{E}_H$ **then**
            $\mathcal{T}_{uv}^i \Leftarrow \mathcal{T}_{uv}^i \backslash \{c\}$;
        if $\mathcal{T}_{uv}^i = \emptyset$ **then**
          $A \Leftarrow \{j | (l_{uv}^i, l_{pq}^j) \in \mathcal{E}_H, \forall p, q; \mathrm{Win}(j) = \mathrm{TRUE}\}$;
          $\phi^*(i) \Leftarrow \min(\phi^*(i), \min_{j \in A} \phi(j))$;
      $p(i) \Leftarrow \phi^*(i) \times |I_s(i)|$;

allocation rule and finding out a critical bidder for each winning agent, is directly based on Theorem 1.

V. A Truthful Auction for Approximately Maximizing Social Welfare

The greedy auction just presented in Sec. IV, while simple and truthful, attempts to maximize social welfare in a heuristic manner, without providing any guarantee. In this section, we set out to design a more sophisticated, randomized, linear programming based truthful auction with proven bounds on approximate social welfare maximization.

We first obtain the linear programming relaxation (LPR) of IP (5), and solve the LPR optimally. We can turn the optimal LPR solution, scaled down by Λ, to a convex combination of integer solutions, through a primal-dual linear program. Here Λ is the upper bound of the integrality gap of IP (5) and the LPR. Viewing this convex combination as specifying a probability distribution over the integer solutions, we arrive at a randomized, truthful in expectation auction with the VCG payment scheme. We prove that it approximately maximizes social welfare.

A. Decomposing the Fractional Solution

Optimally and efficiently solving IP (5) is infeasible, so we resort to an approximation approach. First, we obtain an LPR by allowing the integer variables $f^i_{uv},\ x(c, l^i_{uv})$ to take fractional values in $[0, 1]$. We use $S(\mathbf{w})$ to denote the objective function of the LPR.

Let $\mathbf{f}^{(*)}_{ds}$ be the optimal flow vector from solving the LPR, which also contains agents' winning/losing information. Assume the integrality gap between IP (5) and its LPR is at most $\Lambda \geq 1$, and there is a Λ-approximation algorithm that "verifies" this gap, *i.e.*, it achieves at least $\frac{1}{\Lambda}S(\mathbf{w})$, for any set of $\mathbf{w}$. We employ the technique due to Lavi and Swamy [7]. We utilize the Λ-approximation algorithm and compute in polynomial time a convex decomposition of $\frac{\mathbf{f}^{(*)}_{ds}}{\Lambda}$ into integer solutions in polynomial size. That is, we will have $\rho(l)$ values such that $\frac{\mathbf{f}^{(*)}_{ds}}{\Lambda} = \sum_{l\in\mathcal{I}} \rho(l) f^i_{ds}(l)$, where $\{\mathbf{f}_{ds}(l)\}_{l\in\mathcal{I}}$ is the set of all integer solutions, $\mathcal{I}$ is its index set and $\rho(l) \geq 0$, $\sum_{l\in\mathcal{I}} \rho(l) = 1$.

Now we can view this decomposition as specifying a probability distribution over the integer solutions. A solution $\mathbf{f}_{ds}(l)$ is selected with probability equal to $\rho(l)$. If we set prices in such a way that the expected prices are VCG prices scaled down by Λ, that will make our randomized auction truthful in expectation. Furthermore, the expected social welfare is exactly the value of the LPR scaled by Λ. We next provide details on the decomposition and then prove the upper bound on the integrality gap.

Let $\mathbf{f}^{(*)}_{ds}$ be an optimal solution to the LPR containing all the f^i_{ds} values. We use $\mathbb{Z}(\mathcal{F}) = \{\mathbf{f}_{ds}(l)\}_{l\in\mathcal{I}}$ to denote the set of all integer solutions to the LPR, where $\mathcal{F}$ denotes the feasible region of the LPR, $\mathcal{I}$ is an index set of the integer solutions. The crux of the method is to decide $\rho(l)$ values such that $\frac{\mathbf{f}^{(*)}_{ds}}{\Lambda} = \sum_{l\in\mathcal{I}} \rho(l) f_{ds}(l)$, where $\rho(l) \geq 0$, $\sum_{l\in\mathcal{I}} \rho(l) = 1$. Now one can view this convex combination as specifying a probability distribution over the integer solutions, where a solution $\mathbf{f}_{ds}(l)$ is selected with probability equal to $\rho(l)$. We solve the LP below to obtain the convex decomposition:

$$\text{minimize } \sum_{l\in\mathcal{I}} \rho(l) \tag{8}$$

subject to

$$\begin{aligned} \sum_{l\in\mathcal{I}} \rho(l) f^i_{ds}(l) = \frac{f^{i(*)}_{ds}}{\Lambda} &\quad \forall i \in \mathcal{N} \\ \sum_{l\in\mathcal{I}} \rho(l) \geq \ 1 & \\ \rho(l) \geq \ 0 &\quad \forall l \in \mathcal{I} \end{aligned}$$

The dual of LP (8) is:

$$\text{maximize } \frac{1}{\Lambda}\sum_{i\in\mathcal{N}} \eta^i f^{i(*)}_{ds} + \lambda \tag{9}$$

subject to

$$\begin{aligned} \sum_{i\in\mathcal{N}} \eta^i f^i_{ds}(l) + \lambda \leq 1 &\quad \forall l \in \mathcal{I} \\ \lambda \geq 0 & \\ \eta^i \text{ unconstrained} &\quad \forall i \in \mathcal{N} \end{aligned}$$

To obtain the decomposition, an algorithm that approximately computes the maximum social welfare is needed. We adapt a joint traffic routing and channel assignment solution framework due to Alicherry *et al.* [15] into our setting, with the objective function being the weighted sum of throughput (social welfare in our problem) and with the number of radios at each node being 1. Given that this algorithm may result in fractional flows, we introduce the following modificaiton: for an SN i whose flow is fractional, among SNs in $\{i\} \cup I_s(i)$ that also have fractional flows, find out the SN with maximum $w_i f^i_{ds}$, set its flow to 1 and others to 0. We denote this modified algorithm as A. We show later that Algorithm A "verifies" the upper bound of the integrality gap between (5) and the LPR.

The primal LP (8) has an exponential number of variables, which will take exponential time to solve with the simplex or interior point method [16]. Hence we resort to its dual (9), which has an exponential number of constraints. If we have a separation oracle, we can apply the ellipsoid method to solve the dual (9) and hence (8) in polynomial time [16]. We show a separation oracle for the dual later, so we can efficiently solve this pair of primal-dual LPs. One can view η^i as a valuation. A potential problem is that the η^i values could be negative, whereas A is only for non-negative valuations. However, one can instead use A with the non-negative valuations $\boldsymbol{\eta}^{(+)}$ given by $\eta^{i(+)} = \max(\eta^i, 0)$, and this yields a separation oracle [7].

Claim 1. *Let $\boldsymbol{\eta} = \{\eta^i\}_{i\in\mathcal{N}}$ be any weight vector. $\eta^{i(+)} = \max(\eta^i, 0)$. Given any integer solution $\tilde{\mathbf{f}}_{ds} \in \mathbb{Z}(\mathcal{F})$, one can obtain $\mathbf{f}_{ds}(l) \in \mathbb{Z}(\mathcal{F})$ such that $\sum_{i\in\mathcal{N}} \eta^i f^i_{ds}(l) = \sum_{i\in\mathcal{N}} \eta^{i(+)} \tilde{f}^i_{ds}$.*

Proof: We first exploit the packing property. That is, if $\mathbf{a}_1 \in \mathbb{Z}(\mathcal{F})$ and $\mathbf{a}_2 \leq \mathbf{a}_1$ is integral then $\mathbf{a}_2 \in \mathbb{Z}(\mathcal{F})$. Now we set $f_{ds}^i(l) = \tilde{f}_{ds}^i$ if $\eta^i \geq 0$ and 0 otherwise. Clearly, $\sum_{i \in \mathcal{N}} \eta^i f_{ds}^i(l) = \sum_{i \in \mathcal{N}} \eta^{i(+)} \tilde{f}_{ds}^i$. Since $f_{ds}^i(l) \leq \tilde{f}_{ds}^i$ is integral, by the packing property $\mathbf{f}_{ds}(l) \in \mathbb{Z}(\mathcal{F})$. ■

Now we are ready to show the following lemma:

Lemma 3. *If the optimal solution to (8) is $\boldsymbol{\rho}^*$, then we have $\sum_{l \in \mathcal{I}} \rho^*(l) = 1$.*

Proof: We show that the optimal value of (9) is 1, and hence the lemma follows by strong LP duality. If we simply set $\lambda = 1$, $\eta^i = 0$ for all $i \in \mathcal{N}$, it provides a feasible solution with value 1. We then prove that the optimal value is at most 1 by way of contradiction. Let $(\boldsymbol{\eta}^{(*)}, \lambda^{(*)})$ denote the optimal solution to (9). Suppose $\frac{1}{\Lambda} \sum_{i \in \mathcal{N}} \eta^{i(*)} f_{ds}^{i(*)} + \lambda^{(*)} > 1$. Using Algorithm A and Claim 1, we can compute a social welfare maximizing feasible solution $\mathbf{f}_{ds}(l)$, such that

$$\begin{aligned}\sum_{i \in \mathcal{N}} \eta^{i(*)} f_{ds}^i(l) &\geq \frac{1}{\Lambda} S(\boldsymbol{\eta}^{(*)(+)}) \\ &= \frac{1}{\Lambda} \sum_{i \in \mathcal{N}} \eta^{i(*)(+)} f_{ds}^{i(*)} \\ &\geq \frac{1}{\Lambda} \sum_{i \in \mathcal{N}} \eta^{i(*)} f_{ds}^{i(*)}. \end{aligned} \tag{10}$$

Now we have $\sum_{i \in \mathcal{N}} \eta^{i(*)} f_{ds}^i(l) + \lambda^{(*)} \geq \frac{1}{\Lambda} \sum_{i \in \mathcal{N}} \eta^{i(*)} f_{ds}^{i(*)} + \lambda^{(*)} > 1$, which contradicts the first set of inequalities of (9), thereby contradicting the feasibility of $(\boldsymbol{\eta}^{(*)}, \lambda^{(*)})$. ■

The above lemma shows that without being more restrictive, the inequality $\frac{1}{\Lambda} \sum_{i \in \mathcal{N}} \eta^i f_{ds}^{i(*)} + \lambda \geq 1$ can be added to the dual (9). We will run the ellipsoid method to solve this dual LP. The first set of inequalities of (9) will be the violated inequalities returned by the separation oracle during the execution of the ellipsoid method. The separation oracle is, at a point $(\boldsymbol{\eta}, \lambda)$, if $\frac{1}{\Lambda} \sum_{i \in \mathcal{N}} \eta^i f_{ds}^{i(*)} + \lambda > 1$, then we can use Algorithm A and Claim 1 to find an $\mathbf{f}_{ds}(l)$ for which the constraints of (9) is violated; otherwise, we use the half space $\frac{1}{\Lambda} \sum_{i \in \mathcal{N}} \eta^i f_{ds}^{i(*)} + \lambda \geq 1$ to cut the current ellipsoid. Since the ellipsoid method is guaranteed to take at most a polynomial number of steps, it will return a set of solutions $\{\mathbf{f}_{ds}(l)\}_{l \in \mathcal{I}}$ that is polynomial in size. Then we can plug back these solutions to (8), leading to a linear program with a polynomial number of variables and constraints, which we can solve to recover $\rho(l)$'s that sum to 1.

B. Studying the Integrality Gap

Now we investigate the integrality gap of the IP (5) and the LPR. In the LPR, fractional channel allocation is directly related to link flows, which can be viewed as the fraction of time a specific link is active. Similarly, we can turn the constraint (2) into the following *Link Scheduling Constraint* [15]:

$$f_{uv}^i + \sum_{l_{pq}^j : (l_{pq}^j, l_{uv}^i) \in \mathcal{E}_H} f_{pq}^j \leq 1. \tag{11}$$

where we assume there is only one channel, for ease of exposition. We will argue later that the upper bound of the integrality gap between IP (5) and the LPR does not change when considering multiple channels.

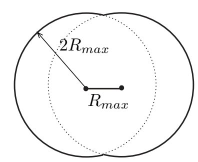

Fig. 3: The circumstance formed by interference regions of one link

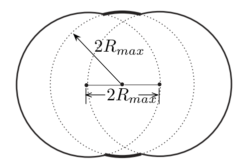

Fig. 4: The circumstance formed by interference regions of two links.

We first introduce the following lemma for a single link, based on geometric arguments:

Lemma 4. *For any channel $c \in \mathcal{C}$,*

$$x(c, l_u^i v) + \sum_{l_{pq}^j : (l_{pq}^j, l_{uv}^i) \in \mathcal{E}_H} x(c, l_p^j q) \leq \alpha(\Delta) \tag{12}$$

where $\alpha(\Delta)$ is a constant that depends only on Δ.

Proof: For a given link l_{uv}^i, we need to find the maximum number of links that interferes with l_{uv}^i yet does not interfere with one another. For simplicity, we prove the following for $\Delta = 2$; proofs for other values of Δ can be derived similarly. The worst case happens when $R_T(u^i) = R_T(v^i) = R_{\max}$ and $d(u^i, v^i) = R_{\max}$. We use $U(u^i, v^i)$ to denote the region formed by the union of two circles with radius $2R_{\max}$ and centre u^i and v^i, respectively. Then the problem becomes placing a max number of points on the circumference of $U(u^i, v^i)$ that are at least $2R_{\max}$ apart. The circumference is composed of two major arcs with length $\frac{4\pi}{3}$, shown in Fig. 3. Among these links, every "independent" set is of size at most 8. Hence in this case $\alpha(\Delta) = 9$. ■

We then investigate the integrality gap based on the worst case analysis of multihop transmissions.

Lemma 5. *For any channel $c \in \mathcal{C}$ and an SN i with a path that has L hops, there are at most $g(L, \Delta)$ interference free SNs among $I_s(i)$ if all of them are assigned channel c, where $g(L, \Delta)$ is a function that only depends on L and Δ.*

Proof: Considering the $\Delta = 2$ case as well, for an end-to-end path, the worst case happens when the path is formed as a straight line, and all the interference-free links along the path belong to different SNs. Adding one link into a path increases the length of circumstance by $4 \cdot \arcsin \frac{1}{4}$, which is due to the bold arcs shown in Fig. 4, where a two-hop path is formed. Then for an SN i with path length L, we have

$$g(L, 2) = \begin{cases} L + 7 & L \text{ is odd} \\ L + 6 & L \text{ is even} \end{cases} \tag{13}$$

We can loosen this bound to $L + 7$, for all L. Similar linear functions of L can be derived with other values of Δ. ■

Theorem 3. *Assume that a SN's path is at most $L_{\max}$-hops. Then the integrality gap between the IP (5) and the LPR is at most $\Lambda = g(L_{\max}, \Delta) + 1$.*

Proof: For an SN i and a single channel c, we know from Lemma 5 that there are at most $g(L, \Delta)$ interference free SNs among $I_s(i)$, ensuring every hop obeys (12). In the worst case, the integral solution picks only one SN from at most $g(L,\Delta)+1$ SNs. Since this is true for any SN and $g(L,\Delta)$ is an increasing function of L, the lemma follows for a single channel case.

If there are $|\mathcal{C}|$ channels, for an SN i, we can imagine that the maximum independent set of a link is duplicated into $|\mathcal{C}|-1$ copies, so that the integral solution will pick SNs from less than $(|\mathcal{C}|-1)g(L,\Delta)L+(|\mathcal{C}|-1)L+1=(|\mathcal{C}|-1)(g(L,\Delta)+1)L$ SNs. Since the integral solution picks at least $(|\mathcal{C}-1|)L+1$ SNs (picks i, and $|\mathcal{C}-1|$ SNs per link along i's path), the integrality gap is at most

$$\begin{aligned}&\frac{(|\mathcal{C}|-1)(g(L_{\max},\Delta)+1)L_{\max}+1}{(|\mathcal{C}|-1)L_{\max}+1}\\ \leq&\frac{(|\mathcal{C}|-1)(g(L_{\max},\Delta)+1)L_{\max}}{(|\mathcal{C}|-1)L_{\max}}\\ \leq&g(L_{\max},\Delta)+1 \qquad (14)\end{aligned}$$

■

Since Algorithm A modifies a fractional flow to 1 among $i\cup I_s(i)$, it apparently "verifies" Λ.

C. A Randomized Approximation Auction

We now present the design of our randomized auction, in Algorithm 3. It first solves the LPR to obtain the optimal $\mathbf{f}_{ds}^{(*)}$. Then a decomposition technique described previously is employed to compute a feasible set of allocations and a probability distribution. Next, a solution is chosen according to its associated probability. Since the ellipsoid method and A both run in polynomial time, our auction is computationally efficient. For each bidder i that wins and achieves throughput $f_{ds}^i(l)$ with probability $\rho(l)$, the payment is calculated as follows:

$$p(i)=\frac{1}{f_{ds}^{i(*)}}(\sum_{j\neq i}b_jy(j)-\sum_{j\neq i}b_jf_{ds}^{i(*)}) \qquad (15)$$

where $\mathbf{y}$ is obtained by recomputing the LPR with $b_i=0$.

Algorithm 3 *A truthful-in-expectation auction*

1. **Input**: Set of channels $\mathcal{C}$, all the compound bids $\mathfrak{B}_i=(G^i(\mathcal{E}^i,\mathcal{V}^i),s^i,d^i,b_i)$, conflict graph $H(\mathcal{E}_H,\mathcal{V}_H)$.
2. Solve the LPR, obtain optimal solution $\mathbf{f}_{ds}^{(*)}$;
3. Use the ellipsoid method and Algorithm A on (9) with $\mathbf{f}_{ds}^{(*)}$, obtain a polynomially sized set of $\{\mathbf{f}_{ds}(l)\}$;
4. Solve (8) with $\mathbf{f}_{ds}^{(*)}$ and $\{\mathbf{f}_{ds}(l)\}$, finding the $\rho(l)$ values;
5. Pick some solution $\mathbf{f}_{ds}(l)$ with probability $\rho(l)$;
6. Let $p(i)\Leftarrow 0$, $\forall i\in\mathcal{N}$;
7. **for all** i such that $f_{ds}^i(l)=1$ **do**
8. Compute $S(b_i=0,\mathbf{b}_{-i})$ with the LPR. Let $\mathbf{y}$ be the solution;
9. $p(i)\Leftarrow\frac{1}{f_{ds}^{i(*)}}(\sum_{j\neq i}b_jy(j)-\sum_{j\neq i}b_jf_{ds}^{i(*)})$;

Theorem 4. *The auction shown in Algorithm 3 is truthful in expectation.*

Proof: Fix i, $\mathbf{b}_{-i}$. Let w_i and w_i' be agent i's bid when being truthful and not truthful, respectively. Let $\mathbf{a}$ and $\mathbf{a}'$ be the solutions to the LPR for bids $(w_i,\mathbf{b}_{-i})$ and $(w_i',\mathbf{b}_{-i})$, respectively. The expected utility of i when bidding truthfully is

$$\begin{aligned}\mathbb{E}[u_i(w_i)]&=\frac{a(i)}{\Lambda}[w_i-\frac{1}{a(i)}(\sum_{j\neq i}b_jy(j)-\sum_{j\neq i}b_ja(j))]\\&=\frac{1}{\Lambda}(w_ia(i)+\sum_{j\neq i}b_ja(j)-\sum_{j\neq i}b_jy(j)) \qquad (16)\end{aligned}$$

Since for every bidder i, $\mathbf{a}$ is optimal for $(w_i,\mathbf{w}_{-i})$. Assuming other bidders bid truthfully yields

$$\begin{aligned}\mathbb{E}[u_i(w_i)]&\geq\frac{1}{\Lambda}(w_ia'(i)+\sum_{j\neq i}b_ja'(j)-\sum_{j\neq i}b_jy(j))\\&=\mathbb{E}[u_i(w_i')] \qquad (17)\end{aligned}$$

■

We now arrive at the following theorem:

Theorem 5. *The auction shown in Algorithm 3 achieves a $\frac{1}{\Lambda}$-approximate maximum social welfare in expectation.*

VI. Simulation Results

Our auctions were evaluated through simulations. We first focus on the heuristic auction. For each SN, we randomly and uniformly distribute some nodes in a 1×1 region. Two nodes are connected if their Euclidean distance is at most 0.05. The largest connected component is used as the connected graph for the corresponding SN. All bids are taken from a uniform distribution in the range of $[40,100]$. Since the auction has already been proven to be truthful and the optimal social welfare is hard to obtain, we evaluate its performance in terms of auction efficiency, which is defined as

$$\vartheta=\frac{\sum_{i\in\mathcal{N}}w_if_{ds}^i}{\sum_{i\in\mathcal{N}}w_i} \qquad (18)$$

We then vary the number of channels in the simulations to study the performance of the auction, where all data are averaged over 100 experiments.

We observe that, in general, as the number of channels increases, the auction efficiency increases as well, which verifies the intuition that the more channels, the higher probability for a bidder to win. First we change the number of bidders (SNs), while fixing $\Delta=4$ and the number of nodes for each SN at 300. From Fig. 5a, we can see that our auction in general effectively exploits the increasing number of channels available. Even in the extremely interfering case where there are 100 SNs in the region, the efficiency increases approximately linearly with the number of channels. Fig. 5b (with 50 bidders and 300 nodes for each SN) and Fig. 5c (with 50 bidders and $\Delta=4$) also show the performance of our auction in terms of the severity of interference, by changing Δ and the size of SNs. We can see that the change of Δ's does not hurt the performance too much. However, large sizes of SNs may increase interference significantly, thereby decreasing the auction efficiency, where the connected graph for an SN with 500 nodes distributed can contain more than 150 nodes.

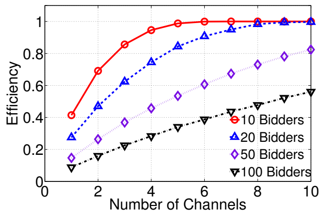

(a) Auction efficiency with different numbers of bidders enrolled.

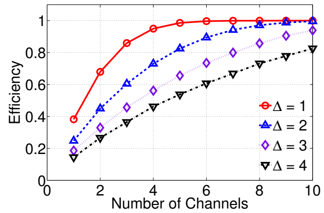

(b) Auction efficiency under different interference situations.

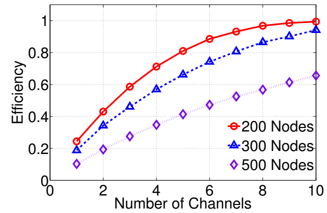

(c) Auction efficiency with different sizes of networks.

Fig. 5: Performance evaluation of the heuristic auction.

We then compare the performance of our auction with two other approaches. One is a greedy auction that only assigns one channel to an SN, in which the same channel cannot be assigned to SNs who interfere with one another. The other one is simply a multi-item auction that greedily assigns channels to links in each SN, without global vision of forming an end-to-end path. We fix $\Delta = 4$, the number of potential nodes for each SN is 450, and the number of bidders is 50. We can see from Fig. 6 that our auction and the single-channel auction perform much better than the multi-item auction. Another observation is that the efficiency of our auction increases faster than the single-channel one as the number of channels increases. This justifies the use of multichannel assignment for each SN.

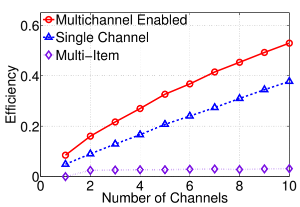

Fig. 6: Comparison of three different auction settings.

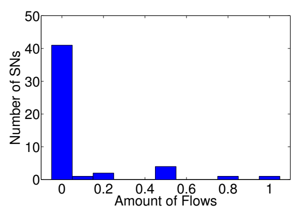

Fig. 7: A histogram of fractional solutions of the LPR.

Fig. 7 shows the distribution of the fractional solution to the LPR for one simulation instance (with 50 bidders, 300 nodes for each SN, one channel and $\Delta = 4$), where the sum of all flows equals 5.67. The agents with relatively large non-zero flows in the solution are shown in Table II. We can see that there is only one agent guaranteed to win and most of them almost always lose (the amount of flows is approximately 0). In the worst case, our randomized auction will select agents with flow amount 1, 0.85, 0.54, and one of agents with flow amount 0.5, and another agent with fractional flow. Note that two agents with flow amount larger than 0.5 must not interfere with each other. Hence, in this experiment, algorithm A can actually achieve $\frac{1}{1.13}$ of the solution to the LPR, hence raising the social welfare with our randomized auction.

TABLE II: Agents with non-zero fractional flows

Agent	2	5	9	12	21	31	35	47
Flow	0.54	0.15	0.46	0.85	0.15	1	0.5	0.5

VII. Conclusion

Secondary spectrum auctions are emerging as a promising approach to efficiently distributing and sharing scarce wireless spectrum. For the first time in the literature, we propose the concept of a *secondary network*, relaxing the over-simplifying assumption on secondary users in existing research. We designed two auctions for spectrum allocation among SNs. The first is a simple, greedy style deterministic auction that heuristically maximizes social welfare. The heuristic auction is truthful due to its monotone allocation rule. The second is a randomized, linear optimization based auction that is not only truthful (in expectation), but also provides proven guarantees on social welfare. In future work, we plan to further improve the performance guarantee of the randomized auction, by proving a tighter bound on social welfare approximation.

References

[1] X. Zhou, S. Gandhi, S. Suri, and H. Zheng, "eBay in the Sky: Strategy-Proof Wireless Spectrum Auctions," in *Proc. ACM MobiCom*, August 2005.
[2] Y. Wu, B. Wang, K. J. R. Liu, and T. C. Clancy, "A Scalable Collusion-Resistant Multi-Winner Cognitive Spectrum Auction Game," *IEEE Trans. Comm.*, vol. 57, no. 12, pp. 3805–3816, 2009.
[3] M. Garey and D. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, 1990.
[4] I. Kash, R. Murty, and D. C. Parkes, "Enabling Spectrum Sharing in Secondary Market Auctions," in *Proc. NetEcon*, June 2011.
[5] A. Gopinathan, Z. Li, and C. Wu, "Strategyproof Auctions for Balancing Social Welfare and Fairness in Secondary Spectrum Markets," in *Proc. IEEE INFOCOM*, April 2011, pp. 3020–3028.
[6] R. Myerson, "Optimal Auction Design," *Mathematics of Operations Research*, vol. 6, no. 1, pp. 58–73, 1981.
[7] R. Lavi and C. Swamy, "Truthful and Near-Optimal Mechanism Design via Linear Programming," in *Proc. IEEE FOCS*, October 2005.
[8] E. H. Clarke, "Multipart Pricing of Public Goods," *Public Choice*, vol. 11, pp. 17–33, 1971.
[9] T. Groves, "Incentives in Teams," *Econometrica: Journal of the Econometric Society*, pp. 617–631, 1973.
[10] W. Vickrey, "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, pp. 8–37, 1961.
[11] J. Huang, R. A. Berry, and M. L. Honig, "Auction Mechanisms for Distributed Spectrum Sharing," in *Proc. Allerton*, September 2004.
[12] M. M. Buddhikot and K. Ryan, "Spectrum Management in Coordinated Dynamic Spectrum Access Based Cellular Networks," in *Proc. IEEE DySPAN*, November 2005, pp. 299–307.
[13] X. Zhou and H. Zheng, "TRUST: A General Framework for Truthful Double Spectrum Auctions," in *Proc. IEEE INFOCOM 2009*, April 2009.
[14] J. Jia, Q. Zhang, Q. Zhang, and M. Liu, "Revenue Generation for Truthful Spectrum Auction in Dynamic Spectrum Access," in *Proc. ACM MobiHoc*, May 2009, pp. 3–12.
[15] M. Alicherry, R. Bhatia, and L. Li, "Joint Channel Assignment and Routing for Throughput Optimization in Multi-radio Wireless Mesh Networks," in *Proc. ACM MobiCom*, August 2005, pp. 58–72.
[16] G. B. Dantzig and M. N. Thapa, *Linear Programming 2: Theory and Extensions*. Springer, 2003.