

Secondary Spectrum Auctions for Symmetric and Submodular Bidders

MARTIN HOEFER, RWTH Aachen University
THOMAS KESSELHEIM, RWTH Aachen University

We study truthful auctions for secondary spectrum usage in wireless networks. In this scenario, n communication requests need to be allocated to k available channels that are subject to interference and noise. We present the first truthful mechanisms for secondary spectrum auctions with symmetric or submodular valuations. Our approach to model interference uses an edge-weighted conflict graph, and our algorithms provide asymptotically almost optimal approximation bounds for conflict graphs with a small inductive independence number $\rho \ll n$. This approach covers a large variety of interference models such as, e.g., the protocol model or the recently popular physical model of interference. For unweighted conflict graphs and symmetric valuations we use LP-rounding to obtain $O(\rho)$ -approximate mechanisms; for weighted conflict graphs we get a factor of $O(\rho \cdot (\log n + \log k))$. For submodular users we combine the convex rounding framework of [Dughmi et al. 2011] with randomized meta-rounding to obtain $O(\rho)$ -approximate mechanisms for matroid-rank-sum valuations; for weighted conflict graphs we can fully drop the dependence on k to get $O(\rho \cdot \log n)$. We conclude with promising initial results for deterministically truthful mechanisms that allow approximation factors based on ρ .

Categories and Subject Descriptors: F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Combinatorial Auctions, Wireless Networks, Radio Interference, Mechanism Design, Spectrum Auctions

1. INTRODUCTION

The development of wireless networks crucially relies on successful management of the frequency spectrum to provide reliable network access. Nowadays, spectrum allocation is static – service providers (so-called *primary users*) can obtain nation-wide licenses for channels in governmental spectrum auctions. This practice is inefficient and problematic: While primary users often use their spectrum bands only in selected local areas, new and innovative applications suffer in their development, because global licenses are difficult to obtain or generally unavailable. A major research effort is currently underway in computer science and engineering to overcome this artificial scarcity and let primary users open their bands in local areas for so-called secondary usage. *Auctions* are attractive to coordinate secondary spectrum usage, as they allow implementing social or monetary goals in a market with self-interested participants having private information. Interest in secondary spectrum auctions has increased significantly in recent years (see [Hoefler et al. 2011; Gopinathan et al. 2011; Gopinathan

This work has been supported by DFG through grant Ho 3831/3-1 and through UMIC Research Centre at RWTH Aachen University.

Author's addresses: M. Hoefler and T. Kesselheim, Department of Computer Science, RWTH Aachen University, {mhoefler, kesselheim}@cs.rwth-aachen.de.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

EC'12, June 4–8, 2012, Valencia, Spain.

Copyright 2012 ACM 978-1-4503-1415-2/12/06...\$10.00.

and Li 2011; Zhou et al. 2008; Zhou and Zheng 2009], and [Berry et al. 2010] for a general discussion), but the algorithmic and strategic problems are still poorly understood.

In secondary spectrum markets, a natural regulatory goal is to maximize social welfare, i.e., the total valuation or benefit of the channel allocation to the secondary users. As constraint for the allocation, the assigned channels must allow successful transmission in the presence of interference and noise. Positioning and interference situation is often known or can sometimes even be observed publicly, but valuations are private information of the users and have to be collected by the algorithm. In this process, secondary users have an obvious incentive to manipulate the algorithm by misreporting their valuation. In this paper, we therefore strive to design *truthful mechanisms* that allocate channels and use payments to motivate users to reveal their values truthfully.

This scenario represents a novel and non-trivial extension of combinatorial auctions. In combinatorial auctions we have to allocate k indivisible items (channels) to n bidders (users). Each bidder v has a valuation $b_v(S)$ for any subset S of items. The goal is to maximize social welfare, i.e., the sum of (reported) valuations for the assigned item sets. Secondary spectrum auctions extend this model by allowing to give a single item/channel to *multiple users* if the set of users is feasible in terms of interference. Interference can be modeled in various ways, and we follow the approach of [Hoefer et al. 2011] where users are vertices in a publicly known edge-weighted conflict graph. A set of users is feasible for a channel if they form an independent set in the graph, for a suitably defined notion of independent set. This approach covers virtually all existing interference models in the literature [Hoefer et al. 2011; Wan et al. 2009]. For instance, if users are communication requests in the physical model of interference, we can use edge weights corresponding to the affectance between requests, and feasibility due to bounded signal-to-interference-plus-noise ratio (SINR) is then equivalent to having an independent set (as defined below, see also [Hoefer et al. 2011]).

Interestingly, conflict graphs resulting from popular interference models (e.g., protocol model [Wan 2009] or physical model [Kesselheim and Vöcking 2010; Hoefer et al. 2011; Kesselheim 2011]) have a small *inductive independence number* ρ . The wide applicability of this non-standard graph parameter for algorithm design is only recently starting to be explored [Akcoglu et al. 2000; Ye and Borodin 2009; Chen et al. 2010]. For our secondary spectrum auctions it allows to bypass well-known lower bounds of $\Omega(n^{1-\epsilon})$ for approximating independent set and derive significantly improved guarantees based on ρ [Hoefer et al. 2011]. However, even in ordinary combinatorial auctions with $\rho = 1$ any efficient algorithm can only achieve a factor of essentially $\min\{n, \sqrt{k}\}$ unless we make additional assumptions on the user valuations [Mirrokni et al. 2008; Lehmann et al. 2002].

1.1. Contribution

In this paper, we design randomized auctions for spectrum markets, where secondary users strive to acquire one or more of a set of available channels. Making additional assumptions on the user valuations allows us to bypass the $\min\{n, \sqrt{k}\}$ lower bound and to significantly improve previous results. We examine the two prominent classes of symmetric and submodular valuations. Both classes occupy a central position in the literature on combination auctions, and they have very natural and intuitive interpretations in the context of secondary spectrum auctions.

Symmetric valuations are the analog to multi-unit auctions, where each valuation only depends on the *number* of channels rather than the exact subset. This is a natural assumption in a secondary spectrum auction of equally sized channels which all offer very similar conditions. Submodularity is economically interpreted as diminish-

ing marginal returns. A common representative are coverage valuations, where users pick elements each covering a certain range, and the value is the total covered area. This is a natural assumption, e.g, when secondary users are transmitters that strive to be received by as many mobile stations as possible, where each of the latter operates on a fixed subset of channels.

For symmetric valuations (see Section 3) we use the intuition of multi-unit auctions and round a suitably defined linear program yielding only an assignment of numbers of channels. Using these numbers an independent set for each channel is then created by a greedy approach. This allows to avoid dependence on k and obtain an approximation factor of $O(\rho)$ for unweighted conflict graphs. Note that this is asymptotically almost optimal under standard complexity assumptions. Theorem 5 in [Hoefer et al. 2011] shows that there is no $\rho/2^{O(\sqrt{\log \rho})}$ -approximation unless $P = NP$. Truthfulness is achieved via combination of our approach with the celebrated randomized meta-rounding framework by Lavi and Swamy [Lavi and Swamy 2005]. For edge-weighted conflict graphs, the construction step of independent sets is significantly more involved. The asymmetry of conflicts inherent in edge-weighted graphs require the use of additional concurrent contention resolution methods to partition the rounded set of requests into feasible independent sets. This approach allows to obtain a factor of $O(\rho \cdot (\log n + \log k))$. Our resulting mechanisms are randomized, run in polynomial time, and yield truthfulness in expectation.

For submodular valuations (see Section 4) we focus on matroid-rank-sum valuations, which encompass the most frequently studied submodular valuations. We design randomized mechanisms that fall into the class of maximum-in-distributional range (MIDR) mechanisms. In particular, our approach is along the lines of the convex rounding technique recently pioneered in [Dughmi 2011; Dughmi et al. 2011] and achieves an approximation factor of $O(\rho)$ for unweighted conflict graphs. Again, this is asymptotically almost optimal under standard complexity assumptions. In contrast to the case of symmetric valuations, we can fully omit the dependence on k and show factors of $O(\rho \cdot \log n)$ even for weighted conflict graphs. Our rounding scheme is similar to the Poisson rounding scheme from [Dughmi et al. 2011]. The main difference and complication is again the need to round each channel to an independent set of users. To achieve this, we round independently for each channel and build the required support of independent sets using a randomized meta-rounding technique. Probably the most technical contribution is showing that this rounding scheme preserves the favorable conditioning properties that allow to apply convex optimization techniques to compute the underlying distribution with sufficient precision in expected polynomial time, even for weighted conflict graphs. Our resulting mechanisms are again randomized and provide truthfulness in expectation.

Finally, we also briefly discuss designing deterministic truthful mechanisms (see Section 5). We present a promising initial result, a monotone greedy $O(\rho \cdot \log n)$ -algorithm for a single channel in unweighted conflict graphs. However, this area remains mostly as an interesting and important avenue for future work.

All proofs missing from this extended abstract will be provided in the full version of this paper.

1.2. Related Work

Our paper is connected to recent works designing truthful mechanisms in secondary spectrum markets without [Zhou et al. 2008; Zhou and Zheng 2009] and with non-trivial worst-case approximation guarantees, e.g., for social welfare and fairness [Gopinathan et al. 2011] or revenue [Gopinathan and Li 2011]. However, all these works are restricted to a single channel and unweighted conflict graphs. In contrast,

variants of combinatorial auctions have been studied where items are sold in interger amounts [Bartal et al. 2003] but without restrictions on the set of bidders that receive a copy. To this date the only general (analytical) treatment of approximation algorithms and truthful mechanisms for multi-channel secondary spectrum auctions is [Hoefer et al. 2011] where truthful-in-expectation mechanisms for general user valuations are designed using the inductive independence number in edge-weighted conflict graphs. For unweighted conflict graphs the approximation guarantee is $O(\rho \cdot \sqrt{k})$, for edge-weighted conflict graphs $O(\rho \cdot \sqrt{k} \cdot \log n)$. The former result is asymptotically almost optimal in ρ if $k = 1$ [Trevisan 2001] and in k if $\rho = 1$. The latter lower bound is a well-known result in combinatorial auctions [Mirrokni et al. 2008; Lehmann et al. 2002].

In ordinary combinatorial auctions, these strong lower bounds initiated the study of relevant subclasses of valuations, for an overview see, e.g., [Blumrosen and Nisan 2007]. Symmetric valuations essentially pose a knapsack problem of assigning numbers of items to bidders, and a deterministic truthful greedy 2-approximation [Mu'alem and Nisan 2002] was the first benchmark solution. Since then there has been significant progress including, e.g., approximation schemes for single-minded bidders [Briest et al. 2005], k -minded bidders [Dobzinski and Nisan 2010], or monotone valuations [Dobzinski and Dughmi 2009; Vöcking 2012]. In contrast to these works, we must additionally decompose assigned numbers of channels into an independent set for each single channel. Here we rely on rounding linear programs to ensure that such a decomposition exists and can be found in polynomial time.

For submodular valuations, social welfare maximization without truthfulness is essentially solved. Optimal $(1 - 1/e)$ -approximation algorithms exist even for value oracle access [Vondrák 2008], where each valuation b_v is an oracle that we can query to obtain $b_v(S)$ for a single set S in each operation. This factor cannot be improved assuming either polynomial communication [Mirrokni et al. 2008] in the value oracle model or polynomial-time complexity in general [Khot et al. 2008]. For the strategic setting and general submodular valuations, the best factors are $O\left(\frac{\log m}{\log \log m}\right)$ for truthfulness in expectation [Dobzinski et al. 2010], and $O(\log m \log \log m)$ for universal truthfulness [Dobzinski 2007]. Dughmi et al [Dughmi et al. 2011] recently proposed a convex rounding technique to build truthful-in-expectation mechanisms. Their approach yields an optimal $(1 - 1/e)$ -approximation for the class of matroid-rank-sum valuations. It follows the idea of maximal-in-distributional range (MIDR) mechanisms by defining a range of distributions independent of the valuations and a rounding procedure. Both are designed in a way that finding the optimal distribution over the range for the reported valuations becomes a convex program with favorable conditioning properties. Hence, the optimal distribution can be found using suitable convex optimization methods in expected polynomial time. Truthfulness follows using the Vickrey-Clarke-Groves (VCG) payment scheme. Very recently, Dughmi and Vondrak showed that a similar result cannot be obtained for general submodular valuations in the value-oracle model [Dughmi and Vondrák 2011].

Designing (non-truthful) algorithms for independent set problems in conflict graphs has received significant attention recently, especially for graphs based on the physical model of interference with SINR constraints. If each request has a value of 1 for being in the independent set, asymptotically optimal performance bounds for specific transmission power assignments were obtained when requests are located in various classes of metric spaces [Goussevskaia et al. 2009; Fanghänel et al. 2010; Halldórsson and Mitra 2011]. For the problem where powers can be arbitrarily chosen, there is a constant-factor approximation algorithm [Kesselheim 2011].

The inductive independence number is a non-standard graph parameter that is only recently starting to receive increased attention. Up to our knowledge the parameter has first been used in [Akcoglu et al. 2000], and since then has been rediscovered independently a number of times (see, e.g., [Wan 2009]). Ye and Borodin [Ye and Borodin 2009] recently conducted the first study addressing general issues that arise when using the measure for solving algorithmic problems in unweighted graphs. The eminent usefulness of the parameter for analyzing interference models and spectrum markets was highlighted in [Hoefer et al. 2011].

2. PRELIMINARIES

2.1. Channel Allocation in Spectrum Markets

In secondary spectrum markets there is a set $[k]$ of k available *channels* and a set V of n *users* or *bidders*. Each user $v \in V$ has a *valuation* or *benefit* $b_v : 2^{[k]} \rightarrow \mathbb{R}^+$. A valuation function b_v is called *symmetric* if $b_v(T) = b_v(|T|)$ for all $T \subseteq [k]$. It is *submodular* if $b_v(T \cup T') + b_v(T \cap T') \leq b_v(T) + b_v(T')$ for all $T \subseteq T'$. For submodular valuations we also assume they are *monotone* with $b_v(T) \leq b_v(T')$ for $T \subseteq T'$. A valuation b_v is a *matroid rank sum (MRS)* function if there exists a family of matroid rank functions $u_1, \dots, u_\kappa : 2^{[k]} \rightarrow \mathbb{N}$, and associated non-negative weights $w_1, \dots, w_\kappa \in \mathbb{R}^+$, such that $b_v(T) = \sum_{\ell=1}^{\kappa} w_\ell u_\ell(T)$ for all $T \subseteq [k]$.

To model interference we represent users as vertices in a complete edge-weighted and directed *conflict graph* $G = (V, E, w)$. The weight $w(u, v)$ of edge (u, v) represents the interference that user u creates for user v if both are assigned to the same channel. Interference between users is similar on each channel. A set of users $U \subseteq V$ is *feasible* or an *independent set* if $\sum_{u \in U} w(u, v) < 1$ for all $v \in U$. In *unweighted* conflict graphs all weights $w(u, v) \in \{0, 1\}$ and our definition of independent set is the same as in the classical sense. For many standard interference models, we can define weighted conflict graphs such that independent sets are exactly the sets for which we can have successful simultaneous transmission in the interference model. For instance, the protocol model results in unweighted conflict graphs, or the physical model of interference yields weighted conflict graphs where independent sets are feasible with respect to the SINR; for details see [Hoefer et al. 2011].

The algorithmic challenge in secondary spectrum markets is the *channel allocation problem*. In an optimal solution S , each user v receives a subset of channels $S_v \subseteq [k]$ such that each channel is given to an independent set in the conflict graph and the *social welfare* $b(S) = \sum_{v \in V} b_v(S_v)$ is maximized. In contrast to ordinary combinatorial auctions, an independent set can include more than one user. Our mechanisms cope with this issue using a structural parameter called inductive independence number. Let us define *symmetric weights* by $\bar{w}(u, v) = w(u, v) + w(v, u)$. Then the *inductive independence number* is the smallest number ρ such that there is an ordering π of the vertices satisfying the following condition: For all $v \in V$ and all independent sets $M \subseteq V$, we let $M_v = M \cap \{u \in V \mid \pi(u) < \pi(v)\}$ and have that $\sum_{u \in M_v} \bar{w}(u, v) \leq \rho$. Hence, ρ is the smallest number such that by picking the best ordering we can bound for any $v \in V$ the incoming weight from any independent set among previous vertices to at most ρ . We assume that ρ and the ordering π of V are given. For many interference models and their resulting conflict graphs we can find in polynomial time small upper bounds on ρ and a corresponding ordering witnessing ρ . For example, in the protocol model $\rho = O(1)$ [Wan 2009] and in the physical model $\rho = O(\log n)$ [Kesselheim and Vöcking 2010] or $\rho = O(1)$ [Kesselheim 2011], depending on power control assumptions. In both cases, π orders users with decreasing or increasing distance between sender and receiver.

2.2. Mechanism Design Basics

To avoid that user v will strategically misreport his valuation, we charge payments p_v and make *truthfulness* a dominant strategy. For each user $v \in V$ we ensure that his *quasi-linear utility* satisfies $b_v(S_v) - p_v(b_v, b_{-v}) \geq b_v(S'(v)) - p_v(b'_v, b_{-v})$, where S and S' are our solutions to the channel allocation problem when v reports the true b_v and some possibly other b'_v , respectively. This can be achieved using classic *Vickrey-Clarke-Groves (VCG) payments* if the allocation problem is always solved optimally.

In contrast, efficient truthful mechanisms cannot compute optimal solutions to intractable problems. For some problems, deterministic mechanisms can achieve only trivial approximation guarantees [Papadimitriou et al. 2008]. The situation is much better if we resort to *randomized mechanisms*, which define a distribution D over the set of solutions \mathcal{S} for the channel allocation problem and output an allocation $S \in \mathcal{S}$ according to D . In this case, we aim for *truthfulness in expectation*, i.e., for every $v \in V$

$$\mathbf{E}_{S \sim D} [b_v(S_v) - p_v(b_v, b_{-v})] \geq \mathbf{E}_{S \sim D'} [b_v(S_v) - p_v(b'_v, b_{-v})] ,$$

where D' is the distribution if v reports b'_v instead of b_v . A general technique to design such mechanisms is *maximal-in-distributional range (MIDR)*. Here we fix a set (the range) of distributions \mathcal{D} over \mathcal{S} , where \mathcal{D} is independent of the valuations b_v . The algorithm receives all reported valuations b_v and optimizes exactly over \mathcal{D} to find $D \in \mathcal{D}$ with maximum expected social welfare. Due to exact optimization over \mathcal{D} , the mechanism can use VCG payments to guarantee truthfulness in expectation. The obvious problem in MIDR is designing the distributional range \mathcal{D} (1) large enough to contain a good approximation for every possible vector of user valuations, and (2) small enough to allow for exact optimization over \mathcal{D} in polynomial time. Our mechanisms in Sections 3 and 4 will all be MIDR mechanisms. In Section 5 we also briefly treat designing greedy mechanisms that are truthful and deterministic.

3. SYMMETRIC VALUATIONS

In this section we consider spectrum auctions with *symmetric valuations* in which $b_v(T) = b_v(|T|)$ for all $v \in V$. We first design approximation algorithms and then turn them into truthful MIDR mechanisms following the framework by Lavi and Swamy [Lavi and Swamy 2005].

Let us first concentrate on designing good approximation algorithms. Our approach is to round the following LP relaxation based on $k \cdot |V|$ variables $x_{v,i} \in \{0, 1\}$ indicating if v gets exactly i channels or not. The relaxation reads

$$\begin{aligned} \text{Max.} \quad & \sum_{v \in V} \sum_{i=1}^k b_v(i) \cdot x_{v,i} \\ \text{s.t.} \quad & \sum_{\substack{u \in V \\ \pi(u) < \pi(v)}} \sum_{i=1}^k i \cdot \bar{w}(u, v) \cdot x_{u,i} \leq \rho \cdot k \quad \text{for all } v \in V \\ & \sum_{i=1}^k x_{v,i} \leq 1 \quad \text{for all } v \in V \\ & x_{v,i} \geq 0 \quad \text{for all } v \in V, i \in [k]. \end{aligned} \tag{1}$$

Note that this relaxation does not describe the problem exactly, as an integral solution to the relaxation might not be feasible for the channel allocation problem. In particular, the relaxation does not specify which user receives which channel, but this information is critical for interference and feasibility of the requests.

ALGORITHM 1: LP-Rounding for Symmetric Valuations and Unweighted Conflict Graphs

- 1 Decompose an optimal solution x to LP (1) into two solutions $x^{(1)}$ and $x^{(2)}$ as follows: Set $x_{v,i}^{(1)} = x_{v,i}$ if $i \leq k/2$ and $x_{v,i}^{(1)} = 0$ otherwise; set $x^{(2)} = x - x^{(1)}$.
 - 2 **for** $l \in \{1, 2\}$ **do**
 - 3 **for** $v \in V$ **in increasing order of** π **values do**
 - 4 with probability $\frac{x_{v,i}^{(l)}}{4\rho}$ set $d_v^{(l)} := i$
 - 5 Let $F_v^{(l)} := \{i \in [k] \mid \text{there is no } u \in \Gamma_\pi(v) \text{ with } i \in S_v^{(l)}\}$
 - 6 $S_v^{(l)} = \begin{cases} \text{arbitrary } M \subseteq F_v^{(l)} \text{ with } |M| = d_v^{(l)} & \text{if } |F_v^{(l)}| \geq d_v^{(l)}, \\ \emptyset & \text{otherwise} \end{cases}$
 - 7 Return the better one of the solutions $S^{(1)}$ and $S^{(2)}$
-

We solve the LP relaxation optimally. The computed fractional solution is then decomposed into two solutions $x^{(1)}$ and $x^{(2)}$, that are rounded separately. Based on such a solution, for each user v a preliminary number of channels $d_v^{(l)}$ is determined at random. The probability is proportional to the fractional variables $x_{v,i}^{(l)}$. Having assigned these numbers of channels, we still have to derive a feasible allocation. In this allocation, each user v either gets $d_v^{(l)}$ channels or none.

3.1. Unweighted Conflict Graphs

In the case of unweighted conflict graphs, we use a simple greedy approach to distribute available channels to users, see Algorithm 1. The expected social welfare of the output will decrease only by a factor of $O(\rho)$ under the fractional optimum, which is asymptotically optimal.

THEOREM 3.1. *Algorithm 1 returns a feasible allocation of social welfare at least $b^*/_{16\rho}$ in expectation.*

3.2. Edge-Weighted Conflict Graphs

Allocating the channels is much more involved in the case of edge-weighted conflict graphs due to the asymmetry of interference constraints. In the unweighted case the simple greedy allocation only has to make sure there are no edges to vertices on the same channel. This is unsuitable now since adding a user might violate constraints at previously added users – even though constraints are satisfied for the currently added user.

Having obtained the $d_v^{(l)}$ values in the described way, we first consider only the incoming weight from users of smaller index like in the unweighted case. If the incoming weight from previous users is too high, i.e., $\sum_{u \in V, \pi(u) < \pi(v)} d_u^{(l)} \cdot \bar{w}(u, v) \geq k/32$, we remove all channels from the user and set $d_v^{(l)} := 0$. However, unlike in the unweighted case, this does not yet guarantee the existence of an allocation. The crucial difference occurs in the last step, where the allocation is derived. This step is performed differently for the two solutions of the decomposition. For the case in which each user was assigned at most $k/8$ channels, the allocation is made in a randomized fashion in Algorithm ALLOCATE(1). For the other case, the allocation is made deterministically in Algorithm ALLOCATE(2). Unlike in the unweighted case, in both cases the resulting allocation will not include all users at a time but only allocate channels to a subset of the originally chosen users.

ALGORITHM 2: LP-Rounding for Symmetric Valuations and Weighted Conflict Graphs

- 1 Decompose an optimal solution x to LP (1) into two solutions $x^{(1)}$ and $x^{(2)}$ as follows: Set $x_{v,i}^{(1)} = x_{v,i}$ if $i \leq k/8$ and $x_{v,i}^{(1)} = 0$ otherwise; set $x^{(2)} = x - x^{(1)}$.
 - 2 **for** $l \in \{1, 2\}$ **do**
 - 3 **for** $v \in V$ **in increasing order of** π **values do**
 - 4 With probability $\frac{x_{v,i}^{(l)}}{64\rho}$ set $d_v^{(l)} := i$
 - 5 Set $d_v^{(l)} := 0$ if $\sum_{u \in V, \pi(u) < \pi(v)} d_u^{(l)} \cdot \bar{w}(u, v) \geq k/32$
 - 6 Run Algorithm ALLOCATE(l) on $d^{(l)}$, let $S^{(l)}$ be the result
 - 7 Return the better one of the solutions $S^{(1)}$ and $S^{(2)}$
-

THEOREM 3.2. *Algorithm 2 returns a feasible allocation of social welfare at least $\Omega(b^*/\rho \cdot (\log n + \log k))$ in expectation.*

In order to show the bound, we will show that both LP solutions are rounded to feasible allocations that are in expectation at most a $O(\rho \cdot (\log n + \log k))$ factor worse than the respective LP solution.

As a first step, we analyze the input given in terms of the number of channels for each user. In particular, we show that an allocation satisfying all of these demands simultaneously would in expectation be at most a $1/128\rho$ factor worse than the fractional solution.

PROPOSITION 3.3. *For $l \in \{1, 2\}$ and the expected social welfare of $d^{(l)}$ we have*

$$\mathbf{E} \left[\sum_{v \in V} b_v(d_v^{(l)}) \right] \geq \frac{1}{128\rho} \cdot \sum_{v \in V} \sum_{i=1}^k b_v(i) \cdot x_{v,i}^{(l)} .$$

In the two following subsections, we consider the two allocation algorithms and show that in either case a feasible allocation of social welfare at least $\Omega(\sum_{v \in V} b_v(d_v^{(l)})/(\log n + \log k))$ is computed.

3.2.1. ALLOCATE(1): Allocation algorithm for “small” sets. From a preliminary selection of numbers of channels Algorithm ALLOCATE(1) generates a feasible allocation in which $d_v \leq k/8$ for each $v \in V$ and $\sum_{u \in V, \pi(u) < \pi(v)} d_u \bar{w}(u, v) < k/32$. The idea is that a number of allocations are computed having the property that each user is considered in exactly one of these allocations. Each allocation is computed by first picking a subset H_t of all users by going through the remaining users in decreasing order of π . If a user is selected, we perform k randomized contention resolution steps. We iterate over the k channels, and for each channel we let each user v independently perform a random experiment. With probability $8d_v/k$ it receives this channel tentatively. If the user received d_v channels it keeps the respective channels in this allocation and is dropped from consideration. All other users are allocated in later rounds. The main argument to show that this yields feasibility and provides the desired bound on the approximation factor relies on a suitable tracking of the successes during the contention resolution process.

3.2.2. ALLOCATE(2): Allocation algorithm for “large” sets. The allocation for the case that $d_v \geq k/8$ or $d_v = 0$ for all $v \in V$ is performed by Algorithm ALLOCATE(2). Here, we iterate starting with $t = 1$. Again, a subset H_t of all users is selected by going through the remaining users in decreasing order of π . If for a user v we have $\sum_{v \in H_t} d_v \cdot \bar{w}(u, v) < k/32$, it is added to H_t . However, in this case the allocation is immediately carried out in

ALGORITHM ALLOCATE(1): Channel allocation for users that require at most $k/8$ channels.

```

1 Set  $V_1 := V$ ,  $t := 1$ , and  $H_1 := \emptyset$ .
2 while  $V_t \neq \emptyset$  do
3   for  $u \in V_t$  in decreasing order of  $\pi$  values do
4     Set  $S_u^t := \emptyset$ .
5     if  $\sum_{v \in H_t} d_v \cdot \bar{w}(u, v) < k/32$  then
6       └ Add  $u$  to  $H_t$  and for each  $j \in [k]$  set  $X_{u,j}$  independently to 1 with probability  $8d_u/k$ .
7   for  $v \in H_t$  do
8     For each  $j \in [k]$  set  $Y_{v,j} = 1$  if  $\sum_{u \neq v} \bar{w}(u, v) \cdot X_{u,j} < 1$ .
9     if  $\sum_{j \in [k]} Y_{v,j} \geq d_v$  then
10      └ set  $S_v^t$  to an arbitrary subset of  $d_v$  channels  $j$  with  $Y_{v,j} = 1$ .
11  Set  $V_{t+1} := \{v \in V_t \mid S_v^t = \emptyset\}$ ,  $H_{t+1} := \emptyset$ , and  $t := t + 1$ .
12 Return the best one of the allocations  $S^1, S^2, S^3, \dots$ 

```

ALGORITHM ALLOCATE(2): Channel allocation for users that require at least $k/8$ channels.

```

1 Set  $V_1 := V$ ,  $t := 1$ , and  $H_1 := \emptyset$ .
2 while  $V_t \neq \emptyset$  do
3   for  $u \in V_t$  in decreasing order of  $\pi$  values do
4     Set  $S_u^t := \emptyset$ .
5     if  $\sum_{v \in H_t} d_v \cdot \bar{w}(u, v) < k/32$  then
6       └ Add  $u$  to  $H_t$  and set  $S_v^t = \{1, \dots, d_v\}$ .
7   Set  $V_{t+1} := \{v \in V_t \mid S_v^t = \emptyset\}$ ,  $H_{t+1} := \emptyset$ , and  $t := t + 1$ .
8 Return the best one of the allocations  $S^1, S^2, S^3, \dots$ 

```

a direct way: Each user that is added to H_t is allocated an arbitrary set of d_v channels, e.g. the first ones. This iteration is repeated with the remaining users that did not get allocated anything until every user $v \in V$ has been allocated d_v channels in one iteration t . Finally, the algorithm picks the best of the allocations computed in any single iteration.

LEMMA 3.4. *For $l \in \{1, 2\}$, ALLOCATE(l) returns a feasible allocation of social welfare $\Omega \left(\frac{\sum_{v \in V} b_v(d_v^{(l)})}{\log n + \log k} \right)$.*

3.3. Mechanism Design

To turn our approximation algorithms into truthful mechanisms, we follow the idea by Lavi and Swamy [Lavi and Swamy 2005] using the randomized meta-rounding technique [Carr and Vempala 2002] to obtain a MIDR mechanism. Our approach is similar to the one for general secondary spectrum auctions [Hoefer et al. 2011], for details see the full version of this paper.

THEOREM 3.5. *There is a truthful mechanism for symmetric valuations that runs in polynomial time and returns a feasible allocation representing $O(\rho)$ -approximation for unweighted and a $O(\rho \cdot (\log n + \log k))$ -approximation for edge-weighted conflict graphs.*

4. MATROID-RANK-SUM VALUATIONS

In this section, we treat the class of so-called *matroid rank sum (MRS) valuations*, in which b_v for each bidder is a weighted sum of matroid rank functions. This covers all

frequently considered submodular valuation functions such as, e.g., coverage functions, matroid weighted-rank functions, and any convex combinations of these.

For ordinary combinatorial auctions, Dughmi et al. [Dughmi et al. 2011] present an MIDR mechanism. The range is given by all solutions to a linear relaxation of the item-allocation problem. Rounding is done via a non-standard randomized rounding scheme called *Poisson rounding* in [Dughmi et al. 2011]. Finding the optimal distribution implies finding the fractional allocation that will achieve best social welfare in expectation in the rounding stage. The Poisson scheme is a *convex rounding scheme*, for which finding the best fractional allocation becomes a convex program with objective function being expected social welfare.

Unfortunately, the Poisson rounding scheme is tailored to fit to ordinary combinatorial auctions. The rounding is performed item-wise – when $x_{i,j}$ is the fractional allocation of item j to bidder i , then j is fully given to i with probability $1 - e^{-x_{v,j}}$. With the remaining probability no bidder receives j . Unlike items, the channels in our case can be given to multiple users, and it takes significantly more effort to build a convex rounding scheme. In the following we present our approach for this case. We follow the conventions in [Dughmi et al. 2011], in particular, with respect to representation of MRS valuations using lottery-value oracles. In particular, we will show the following theorem.

THEOREM 4.1. *There is a truthful mechanism for MRS valuations that runs in expected polynomial time and returns a feasible allocation representing a $O(\rho)$ -approximation for unweighted and a $O(\rho \cdot \log n)$ -approximation for edge-weighted conflict graphs.*

4.1. Defining the Range

We define the distributional range \mathcal{D} in this section and discuss why it is sufficiently large to get good approximations. Our starting point are all fractional solutions x fulfilling the following linear constraints:

$$\sum_{\substack{u \in V \\ \pi(u) < \pi(v)}} \bar{w}(u, v) \cdot x_{u,j} \leq \rho \quad \text{for all } v \in V, j \in [k] \quad (2a)$$

$$0 \leq x_{v,j} \leq 1 \quad \text{for all } v \in V, j \in [k] \quad (2b)$$

ALGORITHM 3: Rounding scheme for a given solution x .

```

1 for  $j \in [k]$  do
2   Draw  $p_j$  uniformly for  $[0, 1]$ 
3   Decompose  $(x_{v,j})_{v \in V}$  such that  $x = \alpha \cdot \sum_l \lambda_l g_l$  and  $\sum_l \lambda_l = 1$ 
4   Let  $l'$  be the minimal  $l$  for which  $\sum_{l < l'} \lambda_l < p_j$ 
5   Allocate  $g_{l'}$  tentatively
6   Remove each  $v \in V$  from solution with a further probability of  $p_{v,j} = \frac{1 - e^{-x_{v,j}/(2\alpha)}}{\frac{x_{v,j}}{\alpha}}$ 

```

For each channel we pick a feasible independent set separately in our rounding scheme Algorithm 3. For each channel j the corresponding fractional solution $x_{\cdot,j}$ is decomposed into polynomially many independent sets using parameter α discussed below. The algorithm selects one of these at random. The decomposition can be computed in polynomial time using randomized meta-rounding [Carr and Vempala 2002; Lavi and Swamy 2005] in combination with an appropriate rounding scheme. Afterwards, each user v is removed from the solution by an independent random experiment

rendering the total probability for v to receive channel j to be exactly $1 - e^{-x_{v,j}/2\alpha}$. Note that $p_{v,j}$ must be a valid probability with $p_{v,j} \in [0, 1]$. Here we observe that since numerator and denominator are both positive, $p_{v,j}$ also is. $p_{v,j} \leq 1$ because $1 - e^{x_{v,j}/(2\alpha)} \leq \frac{x_{v,j}}{2\alpha}$, for any $\alpha \geq 1$. Consequently, the range \mathcal{D} is given by all probability distributions resulting from our rounding scheme applied to fractional solutions of (2a) and (2b).

We have to specify the parameter α , which ensures that the decomposition of $x_{\cdot,j}$ exists. We interpret $x_{\cdot,j}$ as solution to a linear program to maximize $\sum_{v \in V} a_v \cdot x_{v,j}$ subject to the constraints (2a) and (2b) for channel j . This is essentially a linear relaxation for a single channel allocation problem with some valuations a_v . We denote by α the integrality gap of this program with respect to *feasible* independent sets (Note that the constraints (2a) allow integer solutions x that represent infeasible independent sets). For this program we can verify an integrality gap of $\alpha = O(\rho \cdot \log n)$ for feasible independent sets using, e.g., the LP-rounding algorithm for edge-weighted conflict graphs from [Hoefer et al. 2011]. For unweighted conflict graphs, the simpler LP-rounding algorithm from [Hoefer et al. 2011] yields $\alpha = O(\rho)$. Here, the simple greedy algorithm of [Akcoglu et al. 2000] (for details see Section 5 below) can even be shown to yield $\alpha = \rho$.

In either case, this allows to construct a decomposition LP and its dual, which can be solved in polynomial time using the ellipsoid method, where the algorithm acts as separation oracle (for details on this method see [Carr and Vempala 2002; Lavi and Swamy 2005]). Note that α can merely be seen as a parameter that serves to scale a fractional solution x into a region where a decomposition into (feasible) integral solutions exists – independent of any objective function. The reason we interpret it as integrality gap of an optimization problem is that the dual of the decomposition LP allows an approximation algorithm verifying the gap to be used to separate the dual and derive the required decomposition in polynomial time. The reason we do not simply radically overestimate α is that it does play a central role when we discuss the approximation factor of our rounding scheme.

For a given distribution, the expected social welfare of the returned allocation is exactly

$$\sum_{v \in V} \sum_{T \subseteq [k]} b_v(T) \prod_{j \in T} (1 - e^{-x_{v,j}/(2\alpha)}) \prod_{j \notin T} e^{-x_{v,j}/(2\alpha)} . \quad (3)$$

For the case of MRS functions, this function is concave, as we will observe in more detail below. Therefore, the best distribution in the range can be arbitrarily approximated by solving a convex program, maximizing the concave objective (3) subject to linear constraints (2a) and (2b).

As previously mentioned, the size of the range affects approximation factor and tractability. Concerning the approximation factor, we can show that the social welfare of the optimal allocation is at most an $O(\alpha)$ -factor above the expected social welfare of the best distribution in the range.

LEMMA 4.2. *The optimal distribution within the range is $O(\alpha)$ -approximate in expectation when valuations are submodular. Hence, in expectation, the solution of our rounding scheme is a $O(\rho)$ -approximation for unweighted and a $O(\rho \cdot \log n)$ -approximation for edge-weighted conflict graphs.*

4.2. Sampling the MIDR Distribution

The expected social welfare when rounding a fractional solution x is given by (3). Fortunately, this function is concave in terms of x meaning an optimal fractional solution can be approximated arbitrarily well in polynomial time. However, to make the mech-

anism truthful in expectation, we are, in principle, required to solve the given convex program *exactly*.

Since this is not possible, Algorithm 4 devises a way to simulate an exact solution in expected polynomial time. It returns an allocation in which each bidder has exactly the same probability as in Algorithm 3 to get a channel. It requires us to compute δ -estimates – a solution x of the convex program such that $x_{v,j}^* - \delta \leq x_{v,j} \leq x_{v,j}^* + \delta$ for all v, j . To simplify the presentation, we assume that this can be computed in time $\text{poly}(n, k, \log(1/\delta))$. Details on this issue will be presented in the full version of this paper.

ALGORITHM 4: Simulating Algorithm 3 with estimates of the optimal convex-program solution.

```

1 for  $j \in [k]$  do
2   Draw  $p_j$  uniformly from  $[0, 1)$  and let  $r$  be the minimal  $t$  for which  $p_j \geq 1 - 2^{-t+1}$ 
3   Set  $x^0 = 0$ 
4   for  $t = 1, \dots, r$  do
5     Compute  $\delta^t$ -estimate  $x^t$ , where  $\delta^t = 1/(n \cdot 2^{t+1})$ 
6     Let  $y_v^t = \max\{y_v^{t-1}, x_{v,j}^t - \delta^t\}$ 
7     Decompose  $y^r - y^{r-1}$  such that  $y^r = 2\alpha \cdot \sum_l \lambda^{r,l} g^{r,l}$  with  $\sum_l \lambda^{r,l} = 2^{-r}$ 
8     Let  $l'$  be the minimum  $l$  such that  $p_j > 1 - 2^{-r-1} + \sum_{l < l'} \lambda_l$ 
9     Tentatively allocate  $g^{r,l'}$ 
10    Remove each  $v \in V$  from solution with further probability  $p_{v,j} = \frac{2\alpha(e^{-y_v^{t-1}/(2\alpha)} - e^{-y_v^t/(2\alpha)})}{y_v^t - y_v^{t-1}}$ 

```

PROPOSITION 4.3. *The desired decomposition $(g^{r,l}, \lambda^{r,l})_l$ exists and can be computed in polynomial time.*

PROPOSITION 4.4. *For the probability of being removed we have $p_{v,j} \in [0, 1]$.*

PROPOSITION 4.5. *For each user $v \in V$ and each channel $j \in [k]$ the probability to receive j is exactly $1 - e^{-x_{v,j}^*/(2\alpha)}$.*

PROPOSITION 4.6. *Assuming that the δ -estimates can be computed in time $\text{poly}(n, k, \log(1/\delta))$, the expected running time of Algorithm 4 is polynomial in n and k .*

5. DISCUSSION AND OPEN PROBLEMS

While the mechanisms presented in previous sections obtain near-optimal guarantees on social welfare, they have some drawbacks for application in practice. A serious problem are running times – for MRS valuations our mechanism obtains polynomial running time only in expectation. For symmetric valuations, we obtain polynomial worst-case running times, but the convex optimization techniques needed to apply randomized meta-rounding often have prohibitive running times for large practical problem instances. Thus, let us briefly discuss designing fast and simple mechanisms. How can we design a good and simple deterministic mechanism to incentivize truth-telling among bidders?

To our knowledge, there are only two algorithmic approaches to the channel assignment problem that yield approximation guarantees in the order of $O(\rho)$. One approach is rounding of suitably relaxed packing LPs, which turned out to be very successful in this and our previous work [Hoefer et al. 2011]. While pairwise independence can be

used to make these algorithms deterministic, they require randomization to guarantee truthfulness and fail for deterministic truthfulness. The other approach was proposed for the simplest case of a single channel and unweighted conflict graphs, i.e., the maximum weighted independent set problem. It is a simple greedy algorithm due to Akcoglu et al [Akcoglu et al. 2000] which first considers vertices one by one in reverse of the ordering of π . If vertex v is under consideration, its current value is subtracted from the value of each backward neighbor. If the value of a vertex drops to 0 or below before it is under consideration in the ordering, this vertex is removed. Finally, the algorithm makes a second pass over the surviving vertices, this time in forward ordering of π , and greedily adds each vertex to the independent set if possible. It can be shown using a local ratio argument that it provides a ρ -approximation [Ye and Borodin 2009].

It is tempting to believe that this algorithm is monotone and delivers a deterministically truthful mechanism. Unfortunately, there exist simple examples where this is not the case. The problem is that the algorithm makes a second pass over the vertices which introduces non-trivial dependencies among bids and acceptance decisions. Nevertheless, we show how to turn it into a monotone algorithm by spending a $\log n$ factor in the approximation guarantee. This is a promising first step towards designing simple truthful deterministic mechanisms with non-trivial approximation guarantees. In contrast to algorithms using the time-intensive solution of convex optimization problems, such quick and simple greedy rules are much more suitable for application in practice. Providing good and simple mechanisms is a major open direction for future work.

ALGORITHM 5: Monotone $O(\rho \cdot \log n)$ -algorithm for Maximum Weighted Independent Set.

```

1 Sort the set of bids  $B = \{b_v \mid v \in V\}$  in decreasing order, let  $b_i$  be the  $i$ -th highest bid
2 for  $i = 1$  to  $n$  do
3   Let  $V_i = \{v \in V \mid b_v \geq b_i\}$  and  $S_i = \emptyset$ 
4   for  $v \in V_i$  in increasing order of  $\pi$  values do
5     if  $N(v) \cap S_i = \emptyset$ , add  $v$  to  $S_i$ 
6 Output  $S = S_{i^*}$  with  $i^* = \arg \max_i \sum_{v \in S_i} b_v$ 

```

THEOREM 5.1. *Algorithm 5 is deterministic and monotone. The computed solution is a $O(\rho \cdot \log n)$ -approximation for the maximum weight independent set problem.*

This represents a promising first step towards designing simple truthful deterministic mechanisms with non-trivial approximation guarantees. In contrast to algorithms using the time-intensive solution of convex optimization problems, such quick and simple greedy rules are much more suitable for application in practice. In addition, the concept of truthfulness in expectation used in the previous sections has drawbacks, e.g., it is not enough to motivate *risk-aware* bidders to reveal their valuations truthfully. While there are many open problems stemming from our work (e.g., improving the approximation bounds for specific interference models), providing good and simple mechanisms for stronger notions of truthfulness is a challenging and arguably the most interesting avenue for future work.

REFERENCES

- AKCOGLU, K., ASPNES, J., DASGUPTA, B., AND KAO, M.-Y. 2000. Opportunity cost algorithms for combinatorial auctions. *CoRR cs.CE/0010031*.
- BARTAL, Y., GONEN, R., AND NISAN, N. 2003. Incentive compatible multi unit combinatorial auctions. In *Proc. 9th Conf. Theoretical Aspects of Rationality and Knowledge (TARK)*. 72–87.

- BERRY, R., HONIG, M., AND VOHRA, R. 2010. Spectrum markets: Motivation, challenges, and implications. *IEEE Communications Magazine*.
- BLUMROSEN, L. AND NISAN, N. 2007. Combinatorial auctions. In *Algorithmic Game Theory*, N. Nisan, É. Tardos, T. Roughgarden, and V. Vazirani, Eds. Cambridge University Press, Chapter 11.
- BRIEST, P., KRISTA, P., AND VÖCKING, B. 2005. Approximation techniques for utilitarian mechanism design. In *Proc. 37th Symp. Theory of Computing (STOC)*. 39–48.
- CARR, R. AND VEMPALA, S. 2002. Randomized metarounding. *Random Struct. Algorithms* 20, 3, 343–352.
- CHEN, D., FLEISCHER, R., AND LI, J. 2010. Densest-subgraph approximation on intersection graphs. In *Proc. 8th Intl. Workshop Approximation and Online Algorithms (WAOA)*. 83–93.
- DOBZINSKI, S. 2007. Two randomized mechanisms for combinatorial auctions. In *Proc. 10th Intl. Workshop Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*. 89–103.
- DOBZINSKI, S. AND DUGHMI, S. 2009. On the power of randomization in algorithmic mechanism design. In *Proc. 50th Symp. Foundations of Computer Science (FOCS)*. 505–514.
- DOBZINSKI, S., FU, H., AND KLEINBERG, R. 2010. Truthfulness via proxies. *CoRR abs/1011.3232*.
- DOBZINSKI, S. AND NISAN, N. 2010. Mechanisms for multi-unit auctions. *J. Artif. Intell. Res.* 37, 85–98.
- DUGHMI, S. 2011. A truthful randomized mechanism for combinatorial public projects via convex optimization. In *Proc. 12th Conf. Electronic Commerce (EC)*. 263–272.
- DUGHMI, S., ROUGHGARDEN, T., AND YAN, Q. 2011. From convex optimization to randomized mechanisms: Toward optimal combinatorial auctions. In *Proc. 43rd Symp. Theory of Computing (STOC)*. 149–158.
- DUGHMI, S. AND VONDRÁK, J. 2011. Limitations of randomized mechanisms for combinatorial auctions. In *Proc. 52nd Symp. Foundations of Computer Science (FOCS)*. 502–511.
- FANGHÄNEL, A., GEULEN, S., HOEFER, M., AND VÖCKING, B. 2010. Online capacity maximization in wireless networks. In *Proc. 22nd Symp. Parallelism in Algorithms and Architectures (SPAA)*. 92–99.
- GOPINATHAN, A. AND LI, Z. 2011. A prior-free revenue maximizing auction for secondary spectrum access. In *Proc. 30th IEEE Conf. Computer Communications (INFOCOM)*. 86–90.
- GOPINATHAN, A., LI, Z., AND WU, C. 2011. Strategyproof auctions for balancing social welfare and fairness in secondary spectrum markets. In *Proc. 30th IEEE Conf. Computer Communications (INFOCOM)*. 3020–3028.
- GOUSSEVSKAIA, O., HALLDÓRSSON, M., WATTENHOFER, R., AND WELZL, E. 2009. Capacity of arbitrary wireless networks. In *Proc. 28th IEEE Conf. Computer Communications (INFOCOM)*. 1872–1880.
- HALLDÓRSSON, M. AND MITRA, P. 2011. Wireless capacity with oblivious power in general metrics. In *Proc. 22nd Symp. Discrete Algorithms (SODA)*. 1538–1548.
- HOEFER, M., KESSELHEIM, T., AND VÖCKING, B. 2011. Approximation algorithms for secondary spectrum auctions. In *Proc. 23rd Symp. Parallelism in Algorithms and Architectures (SPAA)*. 177–186.
- KESSELHEIM, T. 2011. A constant-factor approximation for wireless capacity maximization with power control in the SINR model. In *Proc. 22nd Symp. Discrete Algorithms (SODA)*. 1549–1559.
- KESSELHEIM, T. AND VÖCKING, B. 2010. Distributed contention resolution in wireless networks. In *Proc. 24th Intl. Symp. Distributed Computing (DISC)*. 163–178.
- KHOT, S., LIPTON, R., MARKAKIS, E., AND MEHTA, A. 2008. Inapproximability results for combinatorial auctions with submodular utility functions. *Algorithmica* 52, 1, 3–18.
- LAVI, R. AND SWAMY, C. 2005. Truthful and near-optimal mechanism design via linear programming. In *Proc. 46th Symp. Foundations of Computer Science (FOCS)*. 595–604.
- LEHMANN, D., O’CALLAGHAN, L., AND SHOHAM, Y. 2002. Truth revelation in approximately efficient combinatorial auctions. *J. ACM* 49, 5.
- MIRROKNI, V., SCHAPIRA, M., AND VONDRÁK, J. 2008. Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. In *Proc. 9th Conf. Electronic Commerce (EC)*. 70–77.
- MU’ALEM, A. AND NISAN, N. 2002. Truthful approximation mechanisms for restricted combinatorial auctions. In *Proc. 18th Conf. Artificial Intelligence (AAAI)*. 379–384.
- PAPADIMITRIOU, C., SCHAPIRA, M., AND SINGER, Y. 2008. On the hardness of being truthful. In *Proc. 49th Symp. Foundations of Computer Science (FOCS)*. 250–259.
- TREVISAN, L. 2001. Non-approximability results for optimization problems on bounded degree instances. In *Proc. 33rd Symp. Theory of Computing (STOC)*. 453–461.
- VÖCKING, B. 2012. A universally-truthful approximation scheme for multi-unit auctions. In *Proc. 23rd Symp. Discrete Algorithms (SODA)*. 846–855.
- VONDRÁK, J. 2008. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proc. 40th Symp. Theory of Computing (STOC)*. 67–74.

- WAN, P.-J. 2009. Multiflows in multihop wireless networks. In *Proc. 10th Symp. Mobile Ad Hoc Networking and Computing (MobiHoc)*. 85–94.
- WAN, P.-J., JIA, X., AND YAO, F. F. 2009. Maximum independent set of links under physical interference model. In *Proc. 4th Intl. Conf. Wireless Algorithms, Systems, Applications (WASA)*. 169–178.
- YE, Y. AND BORODIN, A. 2009. Elimination graphs. In *Proc. 36th Intl. Coll. Automata, Languages and Programming (ICALP)*. 774–785.
- ZHOU, X., GANDHI, S., SURI, S., AND ZHENG, H. 2008. eBay in the Sky: Strategy-proof wireless spectrum auctions. In *Proc. 14th Intl. Conf. Mobile Computing and Networking (MOBICOM)*. 2–13.
- ZHOU, X. AND ZHENG, H. 2009. TRUST: A general framework for truthful double spectrum auctions. In *Proc. 28th IEEE Conf. Computer Communications (INFOCOM)*. 999–1007.