A

Approximation Algorithms for Secondary Spectrum Auctions

MARTIN HOEFER, Max-Planck-Institut für Informatik and Saarland University
THOMAS KESSELHEIM, Cornell University
BERTHOLD VÖCKING, RWTH Aachen University

We study combinatorial auctions for secondary spectrum markets, where short-term communication licenses are sold to wireless nodes. Channels can be assigned to multiple bidders according to interference constraints captured by a conflict graph. We suggest a novel approach to such combinatorial auctions using a graph parameter called inductive independence number. We achieve good approximation results by showing that interference constraints for wireless networks imply a bounded inductive independence number. For example, in the physical model the factor becomes $O(\sqrt{k}\log^2 n)$ for n bidders and k channels. Our algorithms can be turned into incentive compatible mechanisms for bidders with arbitrary valuations.

Categories and Subject Descriptors: F.2.2 [**ANALYSIS OF ALGORITHMS AND PROBLEM COMPLEXITY**]: Nonnumerical Algorithms and Problems — Computations on discrete structures

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Combinatorial Auctions, SINR, Inductive Independence Number

1. INTRODUCTION

A major challenge of today's wireless networks and mobile communication is spectrum management, as devices use common frequency bands that are subject to interference between multiple transmitters in the same area. In fact, spectrum allocation has become one of the key problems that currently limits the growth and evolution of wireless networks. The reason is that, traditionally, frequencies were given away to large service providers in a static way by regulators for entire countries. Examples include FCC auctions in the US or the auctions for UMTS and LTE that took place in Europe. However, demands for services vary at different times and in different areas. Depending on time and place this causes frequency bands licensed for one application to become overloaded. On the other hand, different bands are idle at the same time. A promising solution to this problem is to use market approaches that result in a flexible and thus more efficient redistribution of access rights – thereby overcoming the artificial shortage of available spectrum. In this case, parts of the spectrum that are currently unused by so-called *primary users* for the originally intended purpose (such as TV or telecommunication) can be offered to so-called *secondary users*. Licenses for such secondary usage are valid only for a local area.

A sustainable approach (concisely termed "eBay in the Sky" in [Zhou et al. 2008]) to automatically run such a secondary spectrum market is to auction licenses for secondary users

This work appeared as extended abstract in Proc. 23rd Symp. Parallelism in Algorithms and Architectures (SPAA 2011) [Hoefer et al. 2011]. It has been supported by DFG through grant Ho 3831/3-1, UMIC Research Centre at RWTH Aachen University, and Cluster of Excellence MMCI at Saarland University.
M. Hoefer, Max-Planck-Institut für Informatik, Germany, mhoefer@mpi-inf.mpg.de. T. Kesselheim, Department of Computer Science, Cornell University, USA, kesselheim@cs.rwth-aachen.de. B. Vöcking, Department of Computer Science, RWTH Aachen University, Germany, voecking@cs.rwth-aachen.de.

DOI 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

on a regular basis. In this paper, we propose a general framework and efficient algorithms to implement such a secondary spectrum auction. In our model, there are n secondary users who can bid for bundles of the k wireless channels. Depending on the scenario a user can correspond to a base station that strives to cover a specific area or a pair of devices that want to exchange data (e.g., a base station and a mobile device). In order to account for channel aggregation capabilities of modern devices, users should be able to acquire multiple channels. We allow each user v to have an arbitrary valuation $b_{v,T}$ for each subset T of channels. This level of generality is necessary because of different needs, applications, and hardware abilities of the users, but also because of different locations, spectrum availability, and interference conditions. For instance, the presence of a primary user might allow access to a channel only for a subset of mobile devices located in selected areas. We assume no restrictions on the valuation functions, not even monotonicity.

In this paper, we devise approximation algorithms for spectrum allocation on the secondary market with the objective of maximizing social welfare. We focus on the underlying combinatorial problems and describe interference conflicts by an edge-weighted conflict graph. In unweighted graphs, the vertices represent the bidders and the edges represent conflicts such that the feasible allocations for a channel correspond to the independent sets in the conflict graph. For edge-weighted graphs, we extend the definition of independent set to weighted edges by requiring the sum of all incoming weights to be less than 1. Our formulation covers a large number of binary interference models (such as the protocol model). As we will see, edge weights allow to express more realistic models like the physical model. Here, we can even take the effects of power control into account. Before we state our results, we proceed with a formal description of our model.

1.1. A Combinatorial Framework for Spectrum Assignment

There is a set V of n *users* and a set $[k]$ of k *channels* for wireless transmission. Depending on the application scenario, a user can be interpreted as a transmitter that strives to reach mobile devices in a specific geometric area, or a communication request consisting of a sender-receiver pair located in a metric space. We take a more fundamental approach that captures a large variety of these interpretations. We address the nature of interference by assuming that V is the vertex set of a *conflict graph* $G = (V, E)$. A directed edge $e \in E$ represents a possible conflict, i.e., edge (u, v) indicates that channel access of user u creates a conflict at user v when both try to access the channel simultaneously. To measure the severity of the interference caused by a conflict, we later consider edge weights $w(e) \geq 0$. A conflict graph is termed *unweighted* if $w(e) \in \{0, 1\}$ for all $e \in E$ and *edge-weighted* otherwise. For unweighted conflict graphs, we drop the weights and assume that $e \in E \Leftrightarrow w(e) = 1$. For simplicity of presentation, we first stick to unweighted conflict graphs. In Section 3 below, we present a straightfoward extension to the edge-weighted case.

A user $v \subseteq V$ is successful on channel j if the interference from other users that try to access the channel is not too large. In unweighted conflict graphs, if a set $M \subseteq V$ accesses the channel, user $v \in M$ is successful if there is no user $u \in M$ with $(u, v) \in E$. Hence, a subset of users $M \subseteq V$ can implement simultaneous successful channel access if and only if M is an *independent set* in G.

Users are assumed to have *valuations* for successful channel access. For user v, we denote by $b_{v,T} \geq 0$ the valuation when he has successful access to channel set $T \subseteq [k]$. The main goal of this paper is to find allocations of users to channels that maximize the valuation of successful transmissions. More formally, we try to solve the following combinatorial auction problem with conflict graph.

PROBLEM 1 (COMBINATORIAL AUCTION WITH CONFLICT GRAPH). *Given a graph $G = (V, E)$, a natural number k, and a valuation function $b\colon V \times 2^{[k]} \to \mathbb{N}$, find a feasible allocation $S\colon V \to 2^{[k]}$ that maximizes the social welfare $b(S) := \sum_{v \in V} b_{v,S(v)}$.*

An allocation S is called feasible if for all channels $j \in [k]$, the set of vertices that are assigned to this channel, i. e. $\{v \in V \mid j \in S(v)\}$, is an independent set.

Observe that this problem generalizes combinatorial auctions, in which we have to allocate a set of items to a set of users to maximize total valuation of the allocation. In our case, a channel can be thought of as an item that can be given to multiple bidders representing an independent set in the conflict graphs. To embed ordinary combinatorial auctions in our framework, we can assume the conflict graph to be an unweighted clique. Then each independent set consists only of at most one bidder, which implies that each channel/item can be given to at most one bidder. Instead, for general unweighted conflict graphs and $k = 1$, we obtain the standard maximum weight independent set problem.

By relying on the abstraction to conflict graphs, we can model channel assignment in a large variety of interference models, including, e.g., disk graphs or protocol models. Using an appropriate extension of independent set to edge-weighted conflict graphs, we can even treat variants of the physical model of interference. We elaborate on this issue in Sections 3 and 4 below. However, in general graphs the maximum independent set problem is extremely hard to approximate. Fortunately, conflict graphs resulting from prominent interference models have more structure, which we capture using the following graph parameter. For simplicity, we again define the parameter only for unweighted conflict graphs. For the extension to weighted conflict graphs see Section 3 below.

Definition 1.1 (*inductive independence number ρ*). For a graph $G = (V, E)$, the inductive independence number ρ is the smallest number such that there is an ordering π of the vertices satisfying: For all $v \in V$ and all independent sets $M \subseteq V$, we have $|M \cap \{u \in V \mid \{u, v\} \in E, \pi(u) < \pi(v)\}| \leq \rho$.

In words, for every vertex $v \in V$, the size of an independent set in the *backward neighborhood* of v, i.e., the set of neighbors u of v with $\pi(u) < \pi(v)$, is at most ρ. As we discuss in Section 4, conflict graphs derived from various simple models of wireless communication with binary conflicts like, e.g., the protocol model, distance-2 matchings, or disk graphs, have $\rho = O(1)$. The corresponding ordering π is efficiently computable in these cases. We exploit this property in our algorithms. Note that a bounded inductive independence number is significantly less stringent than bounded independence, which postulates a bound for the complete neighborhood and can increase up to $n - 1$ even in disk graphs.

1.2. Our contribution

We devise the first approximation algorithms for the combinatorial auction problem with conflict graph. Our approach is based on a novel LP formulation for the independent set problem. Our main results concern the so-called physical model [Gupta and Kumar 2000] which is common in the engineering community and has also become subject to theoretical work. In binary models of wireless communication usually studied in theoretical computer science, we make the oversimplifying assumption that interference caused by a signal stops at some boundary around the sender, and receivers beyond this boundary are not disturbed by this signal. In contrast, the physical model takes into account realistic propagation effects and additivity of signals. Feasibility of simultaneous transmissions is modeled in terms of signal to interference plus noise ratio (SINR) constraints. We study two variants of this model, one in which signals are sent at given powers (e.g., uniform) and one where the powers are subject to optimization themselves. We show how to represent SINR constraints for both of these variants in terms of an edge-weighted conflict graph and introduce appropriate notions of "independent set" and "inductive independence number" for edge-weighted graphs.

At first, we observe that the inductive independence number ρ for edge-weighted graphs obtained from the physical model (in both variants) is bounded by $O(\log n)$ and the corresponding ordering is efficiently computable. This enables us to bypass the well-known lower bound of $\Omega(n^{1-\varepsilon})$ [Håstad 1999] on the approximability of independent set in general graphs. In particular, we present an LP relaxation capturing both interference constraints and valuations of users for subsets of channels. Similar to ordinary combinatorial auctions, the LP might require an exponential number of valuations $b_{v,T}$ to be written down explicitly. However, we show how to solve the LP using only oracle access to bidder valuations. Our LP based framework is able to handle edge-weighted conflict graphs resulting from the physical model. By rounding the LP optimum, our algorithm achieves an $O(\rho \cdot \sqrt{k} \log n)$ approximation guarantee. Combining this with the bound on ρ gives an $O(\sqrt{k} \log^2 n)$-approximation of the social welfare for spectrum auctions in the physical model (in both variants).

For more simple binary models of wireless communication such as the protocol model, our approach yields an $\mathcal{O}(\rho \cdot \sqrt{k})$-approximation. Using the bounds on ρ mentioned above, this yields an $\mathcal{O}(\sqrt{k})$ approximation guarantee. In this case, we also provide some complementing hardness results. In general, it is hard to approximate the combinatorial auction problem with conflict graphs to a factor of $\mathcal{O}(\rho^{1-\epsilon})$ and to a factor of $O(k^{\frac{1}{2}-\epsilon})$ for any constant $\epsilon > 0$. While for some specific models better approximations exist, in general the bounds provided by our algorithms for binary models cannot be improved in terms of a single parameter ρ or k. In addition, we provide stronger lower bounds for the case of asymmetric channels, in which the conflict graph can be different for each channel. In this case, our algorithm guarantees a factor of $O(\rho \cdot k)$, which is best possible in general.

Our approach can be used to derive incentive compatible mechanisms using the LP-based randomized meta-rounding framework of Lavi and Swamy [2011] for general bidders with demand oracles. In fact, we slightly extend this framework by starting with an infeasible rather than feasible ILP formulation. The approximation algorithm computes a linear combination of feasible solutions approximating the optimal solution of the corresponding LP and then chooses one of these solutions at random. The obtained mechanism is truthful in expectation, i.e., in expectation over the internal randomization of the mechanism each bidder obtains maximum expected utility by reporting the true valuation, no matter what the reported valuations of other bidders are.

Outline. For technical reasons, we present our results in a different order than stated above. We first introduce the basic approach in the context of unweighted conflict graphs in Section 2. The extensions to edge-weighted graphs including a formal definition of inductive independence number are given in Section 3. The aforementioned wireless models (especially the variants of the physical model) are formally introduced in Section 4, where we also show the bounds on the inductive independence number. The results on asymmetric channels are presented in Section 5. Finally, we conclude with recent progress and open problems in Section 6. The application of the randomized meta-rounding technique to obtain truthful mechanisms is discussed in Appendix A.

1.3. Related Work

The idea of establishing secondary spectrum markets has attracted much attention among researchers in applied networking and engineering communities [Zhou et al. 2008; Gandhi et al. 2008; Buddhikot et al. 2005; Ileri et al. 2005]. There are many different fundamental regulatory questions that need to be addressed when implementing such a market. For example it has to be clarified who runs the market and who is allowed to sell and buy spectrum there. Possible actors could be network providers, brokers, regulators and end-users. In addition, it has to be guaranteed that existing services are not harmed. In most of the literature on spectrum markets the technological aspects dominate. Many results in this area are only of qualitative nature, only a few examples (such as [Zhou et al. 2008; Zhou and Zheng 2009; Kash et al. 2011]) do explicitly consider truthfulness. In terms of non-trivial worst-case guarantees on the efficiency of the allocation there is some recent work on single-parameter auctions with unweighted conflict graphs [Gopinathan et al. 2011; Zhu et al. 2012].

Using our combinatorial models based on (edge-weighted) conflict graphs taking into account the bounded inductive independence number, we put aside technological aspects and focus on the underlying combinatorial and algorithmic questions. To the best of our knowledge there is no previous work on auctions using the general framework of edge-weighted conflict graphs and non-trivial provable worst-case guarantees on the efficiency of the allocation in the multi-parameter case. Subsequent to publication of an extended abstract of this paper in [Hoefer et al. 2011], several special cases and refinements of the framework presented here have been treated. We address this recent progress in Section 6.

In contrast, combinatorial auctions have been a prominent research area in algorithmic game theory over the last decade. A variety of works treats auctions with special valuation functions, such as submodular valuations or ones expressible by specific bidding languages. For an introduction see, e.g., [Nisan et al. 2007, Chapters 11 and 12] or [Cramton et al. 2006]. In addition, designing (non-truthful) approximation algorithms for the allocation problems has found interest, most notably for submodular valuations (e.g., [Vondrák 2008; Feige and Vondrák 2006]). More relevant to our work, however, are results that deal with truthful mechanisms for general valuations. Most notably, Lavi and Swamy [2011] and Dobzinski et al. [2012] derive mechanisms using only demand oracles that achieve an $\sqrt{k}$-approximation with truthfulness in expectation and universal truthfulness, respectively. A deterministic truthful $(k/\sqrt{\log k})$-approximation is obtained by Holzman et al. [2004].

Over the last decades, there has been much research on finding maximum independent sets in the context of interference models for wireless networks. One of the simplest models in this area are disk graphs, which are mostly analyzed using geometric arguments. See [Fishkin 2003; Gräf et al. 1998] for a summary on the results and typical techniques. Christodoulou et al. [2010] study combinatorial auctions for geometric objects. Similar to our approach, they present an LP formulation based on a property in terms of an ordering, the fatness of geometric objects.

The inductive independence number has been used before in [Akcoglu et al. 2002; Ye and Borodin 2012] to approximate independent sets within a factor of ρ with a motivation stemming from chordal graphs. However, these works do not consider multiple channels or wireless communication. As their algorithm is not monotone, it is also not directly applicable to truthful auctions.

Algorithmic aspects of the physical model have become popular in theoretical research recently, particularly the problem of scheduling, i.e., partitioning a given set of requests in a small number of classes such that all requests are successful. New challenges arise since graph-theoretic coloring methods cannot be directly applied. For example, there have been a number of results on how to choose powers for short schedule lengths [Fanghänel et al. 2009; Fanghänel et al. 2011; Halldórsson et al. 2013]. A popular method is fixing powers according to some distance-based scheme. For uniform power assignments, a constant-factor approximation algorithm for the problem of finding a maximum independent set (i.e., a maximum set that may share a single channel) is presented in [Goussevskaia et al. 2009]. Near-optimal bounds for all sublinear schemes have been presented in [Halldórsson and Mitra 2011]. An online version of the problem has been studied in [Fanghänel et al. 2013] presenting tight bounds depending on the difference in lengths of the requests. For arbitrary power schemes, a constant-factor approximation algorithm has been obtained in [Kesselheim 2011] and extended to more general scenarios in [Kesselheim 2012].

2. UNWEIGHTED CONFLICT GRAPHS

2.1. Our LP relaxation

One can get a very intuitive LP formulation for the Weighted Independent Set problem by leaving out the integer constraints from the integer linear program formulation.

$$\begin{aligned} \text{Max.} \ & \sum_{v \in V} b_v x_v & \\ \text{s.t.} \ & x_u + x_v \le 1 & \text{for all } \{u,v\} \in E \\ & 0 \le x_v \le 1 & \text{for all } v \in V \end{aligned}$$

This LP can be used to approximate Independent Set within a factor of $(\bar{d}+1)/2$ [Hochbaum 1983; Kako et al. 2009] where $\bar{d}$ is the average vertex degree. However, even for the case of a clique the integrality gap is $n/2$.

In contrast to this edge-based LP formulation, we here present a different LP based on the *inductive independence number* ρ (recall Definition 1.1). As we will see later, in typical conflict graphs the inductive independence number is constant and the corresponding ordering π can be efficiently calculated. Here we use $\Gamma_\pi(v) = \{u \in V \mid \{u,v\} \in E, \pi(u) < \pi(v)\}$ to denote the backward neighborhood of v. This allows to use the following LP relaxation that has one constraint for each combination of a vertex and a channel and another one for each vertex.

$$\text{Max.} \ \sum_{v \in V} \sum_{T \subseteq [k]} b_{v,T} x_{v,T} \tag{1a}$$

$$\text{s.t.} \ \sum_{u \in \Gamma_\pi(v)} \sum_{\substack{T \subseteq [k] \\ j \in T}} x_{u,T} \le \rho \qquad \text{for all } v \in V,\ j \in [k] \tag{1b}$$

$$\sum_{T \subseteq [k]} x_{v,T} \le 1 \qquad \text{for all } v \in V \tag{1c}$$

$$x_{v,T} \ge 0 \qquad \text{for all } v \in V,\ T \subseteq [k] \tag{1d}$$

This LP works as follows. For each vertex v and each possible set $T \subseteq [k]$ of channels assigned to this vertex, there is one variable $x_{v,T}$. Due to the bounded inductive independence number all feasible allocations correspond to solutions of the LP. However, not all integer solutions of the LP necessarily correspond to feasible channel allocations. Nevertheless, we will show how to compute a feasible allocation from each solution in Theorem 2.3 below.

LEMMA 2.1. *Let S be a feasible allocation and x be defined by $x_{v,T} = 1$ if $S(v) = T$ and 0 otherwise, then x is a feasible LP solution.*

PROOF. Conditions (1c) and (1d) are obviously satisfied. Let us now consider Condition (1b) for some fixed $v \in V$, $j \in [k]$. Set $M := \{u \in V \mid \pi(u) < \pi(v), j \in S(u)\}$. Since M is an independent set, by definition of the inductive independence number, we have $|M \cap \Gamma_\pi(v)| \le \rho$.

On the other hand, we have $\sum_{u \in \Gamma_\pi(v)} \sum_{T \subseteq [k], j \in T} x_{u,T} = |M \cap \Gamma_\pi(v)| \le \rho$. So x is a feasible LP solution. □

As all coefficients are non-negative, this LP has a packing structure. In particular, we can observe the following decomposition property.

OBSERVATION 2.2. *Let x be a feasible solution to the LP, and $x^{(1)}$ be a vector such that $0 \le x^{(1)}_{v,T} \le x_{v,T}$ for all $v \in V$, $T \subseteq [k]$. Then $x^{(1)}$ and $x^{(2)} := x - x^{(1)}$ are feasible LP solutions as well.*

If there are only polynomially many valuations $b_{v,T}$ non-zero, this LP is solvable in polynomial time. In general, due to the number of subsets, the elementary representation of the $b_{v,T}$ values is exponential in k. We can still solve the LP optimally if bidders can be represented by demand oracles.

2.2. Demand Oracles

If there is an arbitrary number of channels, we must define an appropriate way to query the valuation functions of the requests, as an elementary description becomes prohibitively large. A standard way to deal with this issue in the auction literature is the representation by so-called *demand oracles*. To query the demand oracle of bidder v, we assign each channel i a price p_i. Then the oracle delivers his "demand" $S = \arg\max_{T \subseteq [k]} \left(b_{v,T} - \sum_{i \in T} p_i \right)$, i. e., a bundle that maximizes the utility of v given that he pays the sum of prices of channels in the bundle. In ordinary combinatorial auctions such demand oracles can be used to separate the dual of the underlying LP. We here show that such demand oracles can also be used for the solution of our LP (1). Consider the dual given by

$$\text{Min.} \sum_{v \in V} \sum_{j \in [k]} \rho y_{v,j} + \sum_{v \in V} z_v \tag{2a}$$

$$\text{s. t.} \sum_{\substack{u \in V \\ v \in \Gamma_\pi(u)}} \sum_{j \in T} y_{u,j} + z_v \geq b_{v,T} \qquad \text{for all } v \in V,\, T \subseteq [k] \tag{2b}$$

$$y_{v,j} \geq 0 \qquad \text{for all } v \in V,\, j \in T \tag{2c}$$

In contrast to ordinary combinatorial auctions, we cannot use the solution (y, z) directly as the channel prices. Instead, we choose *bidder-specific* channel prices by $p_{v,j} = \sum_{u \in V, v \in \Gamma_\pi(u)} y_{u,j}$. Using this idea we see that the constraints of the dual are indeed equivalent to upper bounds on the utility with bidder-specific channel prices. By obtaining the demand bundle with highest utility for each player, we find a violated constraint or verify that none exists. This allows to separate the dual LP and to solve it efficiently using the ellipsoid method. This way, we get an equivalent primal LP with only polynomially many constraints. The corresponding primal solution has only polynomially many variables with $x^*_{v,T} > 0$.

2.3. Rounding LP Solutions

Having described the LP relaxation, we now analyze Algorithm 1 computing feasible allocations from LP solutions as follows. First, it decomposes the given LP solution to two solutions $x^{(1)}$ and $x^{(2)}$ (line 1). In $x^{(1)}$ all fractional variables $x_{v,T}$ for sets T with $|T| \geq \sqrt{k}$ are set to zero. To get $x^{(2)}$ the exact opposite is performed. From each one, a feasible allocation is computed and the better one is selected at the end. This means, the algorithm either allocates only sets of size at most $\sqrt{k}$ or only of size at least $\sqrt{k}$. The actual computation of the allocation works the same way for both LP solutions. It consists of two major parts: a rounding stage and a conflict-resolution stage. In the rounding stage (lines 3–4), a tentative allocation is generated as follows. For each vertex v the set of allocated channels $S^{(l)}(v)$ is determined independently at random. Each set $T \neq \emptyset$ is taken with probability $x^{(l)}_{v,T}/2\sqrt{k}\rho$ and with the remaining probability the empty set is allocated.

Conflicts can occur when two adjacent vertices share the same channel. In this case, the conflict is resolved (lines 5–8) by allocating the channel to the vertex with smaller index in the π ordering. The other vertex is removed from the solution by being allocated the empty set.

Algorithm 1: LP rounding algorithm for the combinatorial auction problem with unweighted conflict graphs

decompose x into two solutions $x^{(1)}$ and $x^{(2)}$ by $x^{(1)}_{v,T} = x_{v,T}$ if $|T| \leq \sqrt{k}$ and $x^{(1)}_{v,T} = 0$ otherwise. $x^{(2)} = x - x^{(1)}$
for $l \in \{1, 2\}$ **do**
  for $v \in V$ **do** `/* Rounding Stage */`
    with probability $\frac{x^{(l)}_{v,T}}{2\sqrt{k}\rho}$ set $S^{(l)}(v) := T$
  for $v \in V$ **do** `/* Conflict-Resolution Stage */`
    for $u \in V$ *with* $\pi(u) < \pi(v)$ *and* $\{u, v\} \in E$ **do**
      if $S^{(l)}(u) \cap S^{(l)}(v) \neq \emptyset$ **then**
        $S^{(l)}(v) := \emptyset$
return the better one of the solutions $S^{(1)}$ and $S^{(2)}$

THEOREM 2.3. *For any feasible LP solution x^* with value b^*, Algorithm 1 calculates a feasible allocation S of value at least $b^*/8\sqrt{k}\rho$ in expectation.*

PROOF. The allocations $S^{(1)}$ and $S^{(2)}$ are obviously feasible allocations because if $\{u, v\} \in E$, then $S^{(1)}(u) \cap S^{(1)}(v) = \emptyset$ and $S^{(2)}(u) \cap S^{(2)}(v) = \emptyset$. Therefore, the output is also a feasible allocation.

Let us now bound the expected values of solutions $S^{(1)}$ and $S^{(2)}$. Let $l \in \{1, 2\}$ be fixed. Let $X_{v,T}$ be a 0/1 random variable indicating if $S^{(l)}(v)$ is set to T after the rounding stage. Clearly, we have

$$\mathbf{E}\left[X_{v,T}\right] = \frac{x^{(l)}_{v,T}}{2\sqrt{k}\rho} \quad . \tag{3}$$

Let $X'_{v,T}$ be a 0/1 random variable indicating if $S^{(l)}(v)$ is set to T after the conflict-resolution stage. We consider the event that $X'_{v,T} = 0$, given that $X_{v,T} = 1$, i. e. that v is removed in the conflict-resolution stage after having survived the rounding stage.

LEMMA 2.4. *The probability of being removed in the conflict-resolution stage after having survived the rounding stage is at most 1/2, i.e., $\mathbf{Pr}\left[X'_{v,T} = 0 \mid X_{v,T} = 1\right] \leq 1/2$.*

PROOF. The event can only occur if $X_{u,T'} = 1$ for some $u \in V$ with $\pi(u) < \pi(v)$, $\{u, v\} \in E$, and $T \cap T' \neq \emptyset$. In terms of the random variables $X_{u,T}$ this is

$$\sum_{u \in \Gamma_\pi(v)} \sum_{\substack{T' \subseteq [k] \\ T \cap T' \neq \emptyset}} X_{u,T'} \geq 1 \quad .$$

Using this notation we can bound the probability of the event by using Markov's inequality

$$\mathbf{Pr}\left[X'_{v,T} = 0 \mid X_{v,T} = 1\right] \leq \mathbf{Pr}\left[\sum_{u \in \Gamma_\pi(v)} \sum_{\substack{T' \subseteq [k] \\ T \cap T' \neq \emptyset}} X_{u,T'} \geq 1\right] \leq \mathbf{E}\left[\sum_{u \in \Gamma_\pi(v)} \sum_{\substack{T' \subseteq [k] \\ T \cap T' \neq \emptyset}} X_{u,T'}\right] \quad .$$

We will now show separately that this expectation is at most 1/2 for each of the two possible values of l ($l = 1$ or $l = 2$).

Case 1 ($l=1$). We have:

$$\mathbf{E}\left[\sum_{u\in\Gamma_\pi(v)}\sum_{\substack{T'\subseteq[k]\\T\cap T'\neq\emptyset}} X_{u,T'}\right] \quad\leq\quad \mathbf{E}\left[\sum_{j\in T}\sum_{u\in\Gamma_\pi(v)}\sum_{\substack{T'\subseteq[k]\\j\in T'}} X_{u,T'}\right] \enspace .$$

Due to linearity of expectation this is equal to

$$\sum_{j\in T}\sum_{u\in\Gamma_\pi(v)}\sum_{\substack{T'\subseteq[k]\\j\in T'}} \mathbf{E}\left[X_{u,T'}\right] \enspace .$$

Using Equation (3) and the fact that $x^{(1)}$ is an LP solution, this is

$$\sum_{j\in T}\sum_{u\in\Gamma_\pi(v)}\sum_{\substack{T'\subseteq[k]\\j\in T'}} \frac{x^{(1)}_{u,T'}}{2\sqrt{k}\rho} \quad\leq\quad \sum_{j\in T}\frac{1}{2\sqrt{k}} \enspace .$$

Recall that we only have to deal with sets T for which $|T|\leq\sqrt{k}$ in this case. Hence, the expectation is at most $1/2$, and so is the probability that v is removed in the conflict-resolution stage.

Case 2 ($l=2$). In this case, we have $X_{u,T'}>0$ only for sets T' with $|T'|\geq\sqrt{k}$. This yields for all $u\in V$

$$\sum_{\substack{T'\subseteq[k]\\T\cap T'\neq\emptyset}} X_{u,T'} \leq \sum_{\substack{T'\subseteq[k]\\T'\neq\emptyset}} X_{u,T'} = \sum_{\substack{T'\subseteq[k]\\T'\neq\emptyset}}\sum_{j\in T'}\frac{X_{u,T'}}{|T'|} = \sum_{j\in[k]}\sum_{\substack{T'\subseteq[k]\\j\in T'}}\frac{X_{u,T'}}{|T'|} \leq \frac{1}{\sqrt{k}}\sum_{j\in[k]}\sum_{\substack{T'\subseteq[k]\\j\in T'}} X_{u,T'} \enspace .$$

So, we get

$$\mathbf{E}\left[\sum_{u\in\Gamma_\pi(v)}\sum_{\substack{T'\subseteq[k]\\T\cap T'\neq\emptyset}} X_{u,T'}\right] \quad\leq\quad \mathbf{E}\left[\frac{1}{\sqrt{k}}\sum_{j\in[k]}\sum_{u\in\Gamma_\pi(v)}\sum_{\substack{T'\subseteq[k]\\j\in T'}} X_{u,T'}\right] \enspace .$$

Again, we use linearity of expectation, Equation (3) and the fact that $x^{(2)}$ is an LP solution. This gives us

$$\frac{1}{\sqrt{k}}\sum_{j\in[k]}\sum_{u\in\Gamma_\pi(v)}\sum_{\substack{T'\subseteq[k]\\j\in T'}} \frac{x^{(2)}_{u,T'}}{2\sqrt{k}\rho} \quad\leq\quad \frac{1}{\sqrt{k}}\sum_{j\in[k]}\frac{1}{2\sqrt{k}} \quad\leq\quad \frac{1}{2} \enspace .$$

This bounds the probability for the second case.

In both cases we have $\mathbf{Pr}\left[X'_{v,T}=0 \mid X_{v,T}=1\right]\leq 1/2$. □

Using Lemma 2.4 and Equation (3) we get for all $v\in V$ and $T\subseteq[k]$ that $\mathbf{E}\left[X'_{v,T}\right]\geq\frac{x^{(l)}_{v,T}}{4\sqrt{k}\rho}$. This yields that both calculated solutions $S^{(l)}$ for $l\in\{1,2\}$ have expected value

$$\mathbf{E}\left[b(S^{(l)})\right] = \mathbf{E}\left[\sum_{v\in V}\sum_{T\subseteq[k]} b_{v,T}\cdot X'_{v,T}\right] = \sum_{v\in V}\sum_{T\subseteq[k]} b_{v,T}\cdot\mathbf{E}\left[X'_{v,T}\right] \geq \frac{1}{4\sqrt{k}\rho}\sum_{v\in V}\sum_{T\subseteq[k]} b_{v,T}x^{(l)}_{v,T} \enspace .$$

The better one of the two solutions has expected value

$$\mathbf{E}\left[\max\{b(S^{(1)}), b(S^{(2)})\}\right] \geq \frac{1}{2}\left(\mathbf{E}\left[b(S^{(1)})\right] + \mathbf{E}\left[b(S^{(2)})\right]\right)$$
$$\geq \frac{1}{8\sqrt{k}\rho}\sum_{v\in V}\sum_{S\subseteq[k]} b_{v,S}\left(x^{(1)}_{v,S} + x^{(2)}_{v,S}\right) = \frac{1}{8\sqrt{k}\rho}\sum_{v\in V}\sum_{S\subseteq[k]} b_{v,S}x_{v,S} = \frac{b^*}{8\sqrt{k}\rho} . \quad \square$$

2.4. Hardness Results

In this section we provide lower bounds for the approximation ratios of our algorithms. This shows that in general the above results cannot be significantly improved without further restricting the model. Our results are based on the hardness of approximating independent set in bounded-degree graphs [Trevisan 2001] or general graphs [Håstad 1999]. A first result is that the $\mathcal{O}(\rho)$ algorithm for the case $k = 1$ is almost optimal.

THEOREM 2.5. *For $k = 1$ and for each $\rho = \mathcal{O}(\log n)$ there is no $\rho/2^{\mathcal{O}(\sqrt{\log \rho})}$ approximation algorithm unless* P= NP.

PROOF. Such an algorithm could be used to approximate Independent Set in bounded-degree graphs. Given a graph with maximum degree d its inductive independence number ρ is also at most d. Trevisan Trevisan [2001] shows that there is no $d/2^{\mathcal{O}(\sqrt{\log d})}$-approximation algorithm for all $d = \mathcal{O}(\log n)$ unless P= NP. This directly yields the claim. $\square$

As a second result we can also prove the impact of the number of channels k has to be as large as $\sqrt{k}$.

THEOREM 2.6. *Even for $\rho = 1$ there is no $k^{\frac{1}{2}-\varepsilon}$-approximation algorithm unless* ZPP= NP.

Our framework extends general combinatorial auctions with k items, and this is a standard result in the area [Nisan et al. 2007, Chapter 9] derived from the hardness of independent set in general graphs.

In conclusion, our algorithmic results are supported by almost matching lower bounds in each parameter. Without further restricting the graph properties (which means to use additional properties of an interference model) no significantly better approximation guarantees can be achieved in terms of ρ resp. k. However, this does not prove no $O(\rho + \sqrt{k})$ approximation can exist.

3. EDGE-WEIGHTED CONFLICT GRAPHS

In this section we extend conflicts over binary relations (conflict/no-conflict). In wireless communication, we encounter situations that a radio transmission is exposed to interference by a number of devices relatively far away. If there were only a single one of them, interference would be acceptable but their overall interference is too high. For such aggregation aspects we introduce edge-weighted conflict graphs, in which there is a non-negative weight $w(u, v)$ between any pair of vertices $u, v \in V$.

If a subset $I \subseteq V$ of users access a channel j, then user $v \in I$ is successful on j if $\sum_{u\in I} w(u, v) < 1$. Put differently, a subset $I \subseteq V$ of users can implement successful simultaneous access to a channel if they form an *independent set* in G, where the following definition captures our condition of success and extends the standard definition of independent set to edge-weighted conflict graphs.

Definition 3.1. For a set of users V with weight function $w\colon V \times V \to [0, \infty)$, an *independent set* $M \subseteq V$ is such that $\sum_{u\in M} w(u, v) < 1$ for all $v \in M$.

Note that edge weights need not be symmetric. To generalize the definition of the inductive independence number, it is therefore convenient to use the following symmetric edge weights $\bar{w}(u, v) = w(u, v) + w(v, u)$. It allows us to argue over the sum of degrees in Lemma 3.4 below. The approximation guarantee is only increased by a factor of 2 through this step. In practice, it might be more convenient to use a maximum instead of the sum for $\bar{w}$.

Definition 3.2. The *inductive independence number* of an edge-weighted graph G is the minimum number ρ such that there is a total ordering $\pi \colon V \to [n]$ (bijective function) which fulfills for all vertices v and all independent sets $M \subseteq \{u \in V \mid \pi(u) < \pi(v)\}$ the following condition: $\sum_{u \in M} \bar{w}(u, v) \leq \rho$.

Note that in this definition the sets M are only required to be independent sets with respect to the original edge-weights w, not necessarily $\bar{w}$. Therefore, any feasible allocation corresponds to a solution of the following LP relaxation.

$$\text{Max.} \sum_{v \in V} \sum_{T \subseteq [k]} b_{v,T} x_{v,T} \tag{4a}$$

$$\text{s.t.} \sum_{\substack{u \in V \\ \pi(u) < \pi(v)}} \sum_{\substack{T \subseteq [k] \\ j \in T}} \bar{w}(u, v) \cdot x_{u,T} \leq \rho \qquad \text{for all } v \in V,\ j \in [k] \tag{4b}$$

$$\sum_{T \subseteq [k]} x_{v,T} \leq 1 \qquad \text{for all } v \in V \tag{4c}$$

$$x_{v,T} \geq 0 \qquad \text{for all } v \in V,\ T \subseteq [k] \tag{4d}$$

For the case that valuations are only given by demand oracles, this LP can be solved in polynomial time in a similar way as described in Section 2.2.

In case of unweighted conflict graphs, we could ensure feasibility by guaranteeing that no node has neighbors with a smaller index in the π-ordering. For weighted graphs resolving conflicts in only one direction only with respect to the π-ordering is not enough: In particular, even if the edge weights to neighboring vertices on the same channel *with smaller indices* are at most 1, the edge weights to *all* neighboring vertices on the same channel can be more than 1. While a feasible independent set requires to fulfill the latter condition, a feasible integer solution to the LP might only satisfy the former condition, as the conflicts w.r.t. backwards neighbors are the ones represented in the constraints.

To cope with this issue, we increase the scaling by another factor of 2. We use rounding and conflict resolution as previously to ensure that for each vertex v the sum of edge weights to neighboring vertices that have smaller indices and share a channel with v is at most $1/2$. Formally, a *partly-feasible allocation* is an allocation $S \colon V \to 2^{[k]}$ such that

$$\sum_{\substack{u \in V \\ \pi(u) < \pi(v) \\ S(v) \cap S(u) \neq \emptyset}} \bar{w}(u, v) < \frac{1}{2} \qquad \text{for all } v \in V\ . \tag{5}$$

Rounding LP solutions to such partly-feasible allocations can be carried out in a similar way as Algorithm 1. Algorithm 2 decomposes the given LP solution the same way as Algorithm 1. Afterwards, it also performs two stages. In the rounding stage (lines 2–4), again a tentative allocation is determined randomly by considering the LP solution as a probability distribution.

Afterwards, only a partial conflict resolution (lines 5–8) is performed: If for some vertex v the sum of edge weights to neighbors that have lower π values and share a channel exceeds $1/2$, it is removed from the solution (i. e. it is allocated the empty set). Such a partly-feasible solution satisfies Equation (5).

Algorithm 2: LP rounding algorithm for the combinatorial auction problem with weighted conflict graphs

decompose x into two solutions $x^{(1)}$ and $x^{(2)}$ by $x^{(1)}_{v,T} = x_{v,T}$ if $|T| \leq \sqrt{k}$ and $x^{(1)}_{v,S} = 0$ otherwise. $x^{(2)} = x - x^{(1)}$
for $l \in \{1,2\}$ **do**
  for $v \in V$ **do** `/* Rounding Stage */`
    with probability $\frac{x_{v,T}}{4\sqrt{k}\rho}$ set $S^{(l)}(v) := T$
  for $v \in V$ **do** `/* Partial Conflict-Resolution Stage */`
    set $U(v) := \{u \in V \mid \pi(u) < \pi(v), S^{(l)}(v) \cap S^{(l)}(u) \neq \emptyset\}$
    if $\sum_{u \in U(v)} \bar{w}(u,v) \geq \frac{1}{2}$ **then**
      $S^{(l)}(v) := \emptyset$
return the better one of the allocations $S^{(1)}$ and $S^{(2)}$

LEMMA 3.3. *For any feasible LP solution x^* with value b^*, Algorithm 2 calculates a partly-feasible allocation S of value at least $b^*/16\sqrt{k}\rho$ in expectation.*

PROOF. The allocation is partly feasible since both allocations $S^{(1)}$ and $S^{(2)}$ satisfy Condition (5).

For the value of the solution let us again bound the value of the partly-feasible allocations $S^{(1)}$ and $S^{(2)}$. Again, let us fix $l \in \{1,2\}$. Let $X_{v,T}$ be a 0/1 random variable indicating if $S^{(l)}(v)$ is set to T after the rounding stage. This time, we have

$$\mathbf{E}\left[X_{v,T}\right] = \frac{x^{(l)}_{v,T}}{4\sqrt{k}\rho} . \qquad (6)$$

Let $X'_{v,T}$ be a 0/1 random variable indicating if $S^{(l)}(v)$ is set to T after the partial conflict-resolution stage. Again, we consider the event that $X'_{v,T} = 0$, given that $X'_{v,T} = 1$, i. e., that v is removed in the conflict-resolution stage after having survived the rounding stage.

To bound the probability that a vertex is removed in the partial conflict resolution, we observe

$$\mathbf{Pr}\left[X'_{v,T} = 0 \mid X_{v,T} = 1\right] \leq \mathbf{Pr}\left[\sum_{\substack{u \in V \\ \pi(u) < \pi(v)}} \sum_{\substack{T' \subseteq [k] \\ T \cap T' \neq \emptyset}} \bar{w}(u,v) \cdot X_{u,T'} \geq \frac{1}{2}\right]$$

$$\leq 2\mathbf{E}\left[\sum_{\substack{u \in V \\ \pi(u) < \pi(v)}} \sum_{\substack{T' \subseteq [k] \\ T \cap T' \neq \emptyset}} \bar{w}(u,v) \cdot X_{u,T'}\right] = \mathbf{E}\left[\sum_{\substack{u \in V \\ \pi(u) < \pi(v)}} \sum_{\substack{T' \subseteq [k] \\ T \cap T' \neq \emptyset}} \bar{w}(u,v) \cdot 2X_{u,T'}\right]$$

due to Markov's inequality and linearity of expectation. Like in the proof of Lemma 2.4, $\mathbf{E}\left[X_{u,T}\right]$ is again defined in terms of $x^{(l)}$ and $x^{(l)}$ is a feasible LP solution. Therefore, we can apply the arguments used in the proof of Lemma 2.4 in an analogous way and obtain that $\mathbf{Pr}\left[X'_{v,T} = 0 \mid X_{v,T} = 1\right] \leq 1/2$ for both cases $l \in \{1,2\}$.

Algorithm 3: Making a partly-feasible allocation fully feasible

$i := 1$
$V' := V$
while $V' \neq \emptyset$ **do**
  initialize S_i by $S_i(v) := S(v)$ for $v \in V'$ and $S_i(v) := \emptyset$ otherwise
  for $v \in V'$ *in order of decreasing* π *values* **do**
    if $\sum_{u \in V',\, S_i(v) \cap S_i(u) \neq \emptyset} \bar{w}(u,v) < 1$ **then**
      delete v from V' /* v stays in S_i */
    else
      $S_i(v) := \emptyset$ /* v is removed from S_i */
  $i := i + 1$
return the best one of the allocations S_1, S_2, …

Using Equation (6), we get for all $v \in V$, $T \subseteq [k]$ that $\mathbf{Pr}\left[X'_{v,T} = 1\right] \geq \frac{x^{(l)}_{v,T}}{8\sqrt{k}\rho}$. Thus, we can conclude that for $l \in \{1,2\}$, we have $\mathbf{E}\left[b(S^{(l)})\right] \geq \frac{1}{8\sqrt{k}\rho} \sum_{v \in V} \sum_{T \subseteq [k]} b_{v,T} x^{(l)}_{v,T}$.

The expected value of the output is at least

$$\mathbf{E}\left[\max\{b(S^{(1)}), b(S^{(2)})\}\right] \geq \frac{1}{16\sqrt{k}\rho} \sum_{v \in V} \sum_{T \subseteq [k]} b_{v,T} \left(x^{(1)}_{v,T} + x^{(2)}_{v,T}\right) = \frac{b^*}{16\sqrt{k}\rho} \ . \quad \square$$

Given a partly-feasible allocation S, Algorithm 3 implements the necessary additional conflict resolution to derive a fully-feasible one. The algorithm decomposes the partly-feasible allocation to a number of feasible candidate allocations S_1, S_2, …. Each allocation S_i is initialized such that $S_i(v) = S(v)$ if vertex v has been removed from all previous allocations $S_1, \ldots, S_{i-1}$. Otherwise $S_i(v) = \emptyset$. Then a conflict resolution is performed on S_i: The vertices are considered by decreasing indices in the π ordering. If the weight bound is satisfied and the vertex is conflict free, it is assigned to the current allocation S_i and thus removed from the set V' of remaining vertices. Otherwise, if the bound is violated, it stays in the set of remaining vertices V' and is removed from S_i by allocating the empty set. At the end, the best one of the candidate allocations S_i is returned. We will see that each candidate allocation allocates at least half of the remaining vertices a non-empty set. Therefore at most $\lceil \log_2 n \rceil$ candidates are computed and the best one has value at least $b(S)/\lceil \log_2 n \rceil$.

LEMMA 3.4. *Given a (not necessarily feasible) allocation S in which Condition 5 is fulfilled for all $v \in V$, Algorithm 3 calculates a feasible allocation of value at least $b(S)/\lceil \log_2 n \rceil$.*

PROOF. Obviously, by construction all candidates are feasible and so is the output allocation. Next, we prove that we need at most $\lceil \log_2 n \rceil$ iterations of the while loop by showing that in each iteration at most half of the remaining vertices are removed from the allocation. This means the cardinality of V' is at least halved in each iteration. Intuitively, in each iteration even if all surviving nodes are assumed to be in the independent set, for at least half of them the condition for independence is fulfilled with respect to their complete neighborhood. Formally, this will be a result of the more stringent right-hand-sides to the LP-constraints and Markov inequality applied to the degrees of nodes. This shows that we can partition S into $O(\log n)$ feasible independent sets.

More formally, let V'_i be the set V' after the ith iteration of the *while* loop; $V'_0 = V$. Let us fix $i \in \mathbb{N}$, and $v \in V'_{i+1}$. We know that v has been removed from S_i by the algorithm.

Denoting by S'_i the current state of S_i while the algorithm considers v, this only happens if

$$\sum_{\substack{u \in V'_i \\ S'_i(v) \cap S'_i(u) \neq \emptyset}} \bar{w}(u,v) \geq 1$$

Since Equation 5 is obviously also satisfied for S'_i, it has to be

$$\sum_{\substack{u \in V'_i \\ \pi(u) > \pi(v) \\ S'_i(v) \cap S'_i(u) \neq \emptyset}} \bar{w}(u,v) \geq \frac{1}{2} \ .$$

For a vertex $u \in V'_i$ with $\pi(u) > \pi(v)$ we know that the vertex has either been removed from the allocation before (then $u \in V'_{i+1}$) or it stays in S_i (i. e. $S'_i(u) = S_i(u) = S(u)$ and $u \notin V'_{i+1}$). Hence

$$S'_i(u) = \begin{cases} \emptyset & \text{if } u \in V'_{i+1} \\ S(u) & \text{else} \end{cases} \ .$$

Combining these two insights, we get a necessary condition: if $v \in V'_{i+1}$ then

$$\sum_{u \in U_i(v) \setminus U_{i+1}(v)} \bar{w}(u,v) \geq \frac{1}{2} \ ,$$

where $U_i(v) = \{u \in V'_i \mid \pi(u) > \pi(v), S(v) \cap S(u) \neq \emptyset\}$. Summing up all $v \in V'_{i+1}$ we get

$$\sum_{v \in V'_{i+1}} \sum_{u \in U_i(v) \setminus U_{i+1}(v)} \bar{w}(u,v) \geq \frac{1}{2} |V'_{i+1}| \ .$$

On the other hand, we can change the ordering of the sums and use the symmetry of the weights $\bar{w}$ to get

$$\begin{aligned} \sum_{v \in V'_{i+1}} \sum_{u \in U_i(v) \setminus U_{i+1}(v)} \bar{w}(u,v) &= \sum_{u \in V'_i \setminus V'_{i+1}} \sum_{\substack{v \in V'_{i+1} \\ \pi(u) > \pi(v) \\ S(v) \cap S(u) \neq \emptyset}} \bar{w}(u,v) \\ = \sum_{v \in V'_i \setminus V'_{i+1}} \sum_{\substack{u \in V'_{i+1} \\ \pi(u) < \pi(v) \\ S(v) \cap S(u) \neq \emptyset}} \bar{w}(u,v) &< \frac{1}{2} |V'_i \setminus V'_{i+1}| \ , \end{aligned}$$

where the last bound is due to Condition (5).

In combination this yields $|V'_{i+1}| < |V'_i \setminus V'_{i+1}|$, which implies $|V'_{i+1}| < \frac{1}{2}|V'_i|$, meaning less than half of the remaining vertices are removed from the partly-feasible allocation in each iteration. So, since $|V'_0| = n$, we can conclude that $|V'_i| < \frac{1}{2^i} \cdot n$. Therefore, we get $|V'_{\lceil \log_2 n \rceil}| < 1$. Thus the algorithm terminates within $\lceil \log_2 n \rceil$ steps.

By definition, for all vertices $S_i(v) = S(v)$ for exactly one $i \in [\lceil \log_2 n \rceil]$ and $S_i(v) = \emptyset$ else. So $\sum_{i \in [\lceil \log_2 n \rceil]} b(S_i) = b(S)$. This yields for the value of the output

$$\max_{i \in [\lceil \log_2 n \rceil]} b(S_i) \geq \frac{1}{\lceil \log_2 n \rceil} \sum_{i \in [\lceil \log_2 n \rceil]} b(S_i) = \frac{b(S)}{\lceil \log_2 n \rceil} \ . \quad \square$$

As a consequence, the computed feasible allocation has a value that in expectation is at most an $\mathcal{O}(\sqrt{k}\rho \log n)$ factor smaller than that of the optimal LP solution.

4. APPLICATIONS

In the previous sections we have described a general algorithmic approach to channel allocation problems when the underlying conflict graph has bounded inductive independence number. Here we will show that this property is particularly wide-spread among models for interference in wireless communication. Our aim is not to prove optimal bounds in each case but to show why we believe a bounded inductive independence number to be a key insight for understanding algorithmic problems in wireless networking.

The concept of conflict graphs can be applied in two basic scenarios. On the one hand, the task could be to allocate channels to *transmitters*. Each transmitter intends to cover a certain area, e.g., a base station in a cellular network. The interference model defines which transmitters can be assigned the same channels. On the other hand, instead of single transmitters one can consider pairs of network nodes (*links*) that act as sender and receiver. In such a scenario, "users" are no single network nodes but links. Therefore, the vertices of the conflict graph are links, and edges define which links can be assigned the same channels.

The following represents a selected set of prominent models. We discuss further examples in Appendix B.

4.1. Transmitter Scenarios

A very simple, yet instructive model for a transmitter scenario is as follows. We have n transmitters located in the plane at points $p_1, \ldots, p_n \in \mathbb{R}^2$. Each of the transmitters has a transmission range $r_1, \ldots, r_n \in \mathbb{R}_{>0}$. Transmitters may be assigned the same channel if their transmission ranges do not intersect. Under these conditions interference constraints can be modeled by a disk graph. There is an edge between two vertices if the transmission-range disks around the corresponding receivers intersect.

PROPOSITION 4.1. *Disk graphs have an inductive independence number of $\rho \leq 5$.*

PROOF. Let $G = (V, E)$ be a disk graph. Let π be the ordering of vertices by decreasing radius of the corresponding disk. So in other words $V = \{v_1, \ldots, v_n\}$ with $r_1 \geq r_2 \geq \ldots \geq r_n$, where r_i is the radius of the disk around v_i. If the disk representation is given, this ordering can be computed in polynomial time by simply sorting the vertices. Let be $v \in V$ and $M \subseteq V$ be an independent set in G. By definition of the ordering, we have for all $u \in M \cap \Gamma_\pi(v)$ the radius is at least $r_{\pi(v)}$.

To show $|M \cap \Gamma_\pi(v_i)| \leq 5$, we assume $|M \cap \Gamma_\pi(v_i)| \geq 6$. This would yield that there were two vertices v_j and v_k whose angle seen from v_i were at most $60°$. The distance between the centers of v_j and v_k would therefore be no larger than the larger ones of the two distances between the centers of v_i and v_j or v_i and v_k. As $r_j \geq r_i$ and $r_k \geq r_i$, this implies that disks around v_j and v_k intersect. So, there has to be an edge between these two vertices. This contradicts the assumption that M is an independent set and thereby proves the claim. □

4.2. Unweighted Link-Based Scenarios

There are a number of different interference models for link-based scenarios that can be described by some unweighted conflict graph. They are often called graph-based interference models, but to avoid ambiguities we refer to them as *binary interference models*. Due to the large variety, we have to confine ourselves to some selected examples.

Probably the best known binary model is the *Protocol Model* [Gupta and Kumar 2000]. Network nodes are modeled by points located in the plane. A link consisting of sender s and receiver r may be allocated to a channel if and only if for all other senders s' on this channel $d(s', r) \geq (1 + \Delta) d(s, r)$ for some constant $\Delta > 0$.

PROPOSITION 4.2 (WAN [2009]). *For the protocol model, the resulting conflict graph has an inductive independence number of $\rho \leq \lceil \pi / \arcsin(\Delta/2(\Delta+1)) \rceil - 1$.*

The *IEEE 802.11 Model* by Alicherry et al. [2006] is a bidirectional variant of this model, and in this case $\rho \leq 23$ [Wan 2009].

4.3. Physical Model

The models mentioned above go well with graph-theoretic concepts. However, radio transmissions typically decrease asymptotically with increasing distance. The *physical model* captures this property much more accurately and is particularly wide-spread among engineers. Even though the physical model does not fit in the traditional binary graph-theoretic context, it has similar properties allowing it to be expressed using edge-weighted conflict graphs.

In this model, network nodes are located in a metric space. The received signal strength decreases as the distance increases. If a node transmits at a power level p, the signal strength at a distance of d is p/d^α, for a constant α. A transmission is received successfully if ratio of the received signal strength of the intended transmission and the stengths of concurrent transmissions plus ambient noise is above some constant threshold $\beta > 0$. More formally, given pairs of senders s_i and receivers r_i that transmit at power level p_i, receiver r_i can decode the signal from sender s_i successfully if the SINR constraint $\frac{p_i}{d(s_i,r_i)^\alpha} \geq \beta \left(\sum_{j \in M \setminus \{i\}} \frac{p_j}{d(s_j,r_i)^\alpha} + \nu \right)$ is fulfilled. Here M is the set of links transmitting at the same time on the same channel and $\nu \geq 0$ is a constant expressing ambient noise.

Note that we can easily reduce the model to a conflict graph if transmission powers are fixed. Prominent and simple classes of power assignments $p \colon V \to \mathbb{R}_{>0}$ are uniform ($p(v) = 1$) or linear ($p(v) = d(s_v, r_v)^\alpha$) assignments. More generally, we can consider assignments satisfying the following monotonicity constraints. For link ℓ we denote by $d(\ell)$ the distance between sender and receiver of link ℓ. If $d(\ell) \leq d(\ell')$ for two links ℓ, ℓ' then $p(\ell) \leq p(\ell')$ and $\frac{p(\ell)}{d(\ell)^\alpha} \geq \frac{p(\ell')}{d(\ell')^\alpha}$. That is, if the distance between a sender and a receiver is larger, the absolute power is higher (or equal), whereas the received power is smaller (or equal). We furthermore assume that each sender uses slightly more power than would be necessary to reach its receiver in the presence of of only ambient noise but no interference. This is formalized as $\frac{p(\ell)}{d(\ell)^\alpha} \geq 2\beta\nu$.[1]

PROPOSITION 4.3. *The interference constraints in the physical model with fixed transmission power can be represented by a weighted conflict graph. If the power assignment satisfies the above constraints, the resulting inductive independence number is at most $\mathcal{O}(\log n)$.*

PROOF. We choose the edges of the conflict graph to have the following weights. For $\ell = (s, r)$, $\ell' = (s', r')$ we set

$$w(\ell', \ell) = \min \left\{ 1, \frac{\beta}{1+\varepsilon} \cdot \frac{p(\ell')}{d(s', r)^\alpha} \Big/ \left(\frac{p(\ell)}{d(s, r)^\alpha} - \frac{\beta}{1+\varepsilon} \nu \right) \right\} \quad ,$$

$$\text{where } \varepsilon = \frac{\beta}{2} \min_{\ell=(s,r)} \min_{\ell'=(s',r')} \frac{p(\ell)}{d(s', r)^\alpha} \Big/ \frac{p(\ell)}{d(s, r)^\alpha} \quad .$$

By this definition, for any set of links M and any $\ell' = (s', r')$, we have $\sum_{\ell' \in M} w(\ell', \ell) < 1$ is equivalent to $\sum_{\ell'=(s',r') \in M} \frac{p(\ell')}{d(s',r)^\alpha} + \nu < \frac{1+\varepsilon}{\beta} \frac{p(\ell)}{d(s,r)^\alpha}$. This is in turn equivalent to this SINR condition: Any set fulfilling the SINR condition also fulfills this inequality as ε is strictly greater than 0. However, as any term in the sum is larger than $\frac{\varepsilon}{\beta} \frac{p(\ell)}{d(s,r)^\alpha}$, it cannot occur

[1]Note that the constant 2 could be replaced by any other constant strictly larger than 1, see Kesselheim and Vöcking [2010].

that only a strict subset of M is SINR feasible. So a set M fulfills the SINR constraint iff it corresponds to an independent set in the edge-weighted graph.

Note that apart from the factor the edge weights are equal to the notion of *affectance* a_p in [Kesselheim and Vöcking 2010], for which we have the following result. Recall that $d(\ell)$ denotes the distance between sender and receiver of link ℓ.

LEMMA 4.4 (KESSELHEIM AND VÖCKING [2010]). *Let p be a power assignment satisfying Conditions 1 and 2 in [Kesselheim and Vöcking 2010]*

If M is a set of links that can concurrently transmit and ℓ is another link with $d(\ell) \leq d(\ell')$ for all $\ell' \in M$, then $\sum_{\ell' \in M} a_p(\ell', \ell) = \mathcal{O}(1)$ and $\sum_{\ell' \in M} a_p(\ell, \ell') = \mathcal{O}(\log n)$.

This immediately yields the edge-weighted graph to have an inductive independence number $\rho = \mathcal{O}(\log n)$. □

While these bounds hold for general metric spaces, Halldórsson et al. [2013] have recently provided tight bounds of $O(\log \log \Delta)$ on the inductive independence number in the physical model with square-root power assignment and Euclidean distance, where Δ is the ratio of longest and shortest distance among any sender-receiver pair.

Interestingly, we can also use our approach if transmission powers are not given upfront. In this case, our algorithm has to decide about the assignment of links to channels and which transmission powers to use for each link. The first part is solved by LP rounding as above. In the LP we use edge weights ensuring that there is a feasible power assignment for the computed set of links. The second task of power assignment can then be done using a power control procedure by Kesselheim [2011].

Note that, in contrast to the interference models mentioned above, in this case not all feasible solutions (i.e., feasibly scheduled sets of links) correspond to independent sets in the weighted graph. However, for our argument it suffices to observe that each set of feasible links corresponds to an LP solution for some ρ and that integral LP solutions with $\rho = 1$ also correspond to feasible sets of links. Combining these insights with the bounds in [Kesselheim 2011] and the ones we proved above, we obtain the following result.

THEOREM 4.5. *There is a choice of edge weights such that our algorithm in combination with the power control procedure in [Kesselheim 2011] achieves an $\mathcal{O}(\sqrt{k} \log n)$ approximation in fading metrics (see [Halldórsson 2013]) and an $\mathcal{O}(\sqrt{k} \log^2 n)$ approximation in general metrics.*

PROOF. We define the weighted graph as follows. The set of vertices is again the set of all links $\mathcal{R}$. Let π be the ordering from large to small distances between the sender and its receiver. For $\tau = \frac{1}{2 \cdot 3^\alpha \cdot (4\beta + 2)}$, set the weight between two links $\ell = (s, r)$ and $\ell' = (s', r')$ to

$$w(\ell, \ell') = \begin{cases} \frac{1}{\tau} \min\left\{1, \frac{d(s,r)^\alpha}{d(s,r')^\alpha}\right\} + \frac{1}{\tau} \min\left\{1, \frac{d(s,r)^\alpha}{d(s',r)^\alpha}\right\} & \text{if } \pi(\ell) < \pi(\ell') \\ 0 & \text{otherwise} \end{cases} .$$

As $\sum_{\ell' \in \mathcal{L}} w(\ell, \ell') < 1$ implies $\sum_{\ell' = (s', r') \in \mathcal{L}, d(\ell') < d(\ell)} \frac{d(s,r)^\alpha}{d(s,r')^\alpha} + \frac{d(s,r)^\alpha}{d(s',r)^\alpha} \leq \tau$ for all $\ell = (s, r)$ because $\tau \leq 1$, Theorem 3 in [Kesselheim 2011] states that for each independent set in the weighted graph the power-control algorithm calculates a feasible set of links.

On the other hand Theorem 1 in [Kesselheim 2011] shows that under the above edge weights each feasible set of links is also an LP solution for some $\rho = \mathcal{O}(1)$ in fading metrics. Theorem 7 in [Kesselheim 2011] shows $\rho = \mathcal{O}(\log n)$ in general metrics.

In conclusion, this implies that by applying our rounding algorithm to the LP using above defined weights we get a solution, for which we can apply the power assignment of [Kesselheim 2011] to obtain a feasible set of links. The resulting allocation is an $\mathcal{O}(\sqrt{k} \log n)$ approximation for fading metrics and an $\mathcal{O}(\sqrt{k} \log^2 n)$ approximation in general metrics. □

5. ASYMMETRIC CHANNELS

So far, channels were symmetric in terms of interference, which means the same interference model is applied to each channel. In a more general setting, for each of the k channels a different edge set E_j resp. a different edge-weight function w_j for the interference graph is given. Then we have an edge weight function $\bar{w}_j$ for each channel $j \in [k]$. The above LP relaxation can be easily adapted by exchanging $\bar{w}$ by $\bar{w}_j$ in the constraints (1b). In contrast, the analysis of the rounding algorithms internally depends on the assumption of symmetric channels: We used symmetry when reordering the sums in the proof of Lemma 2.4.

However, when exchanging the probability for a vertex v to choose set T by $x^{(l)}_{v,T}/2k\rho$ resp. $x^{(l)}_{v,T}/4k\rho$, the proof of Lemma 2.4 can be carried out the same way without using the symmetry. Hence, for the asymmetric case, we lose a factor of $\mathcal{O}(k \cdot \rho)$ resp. $\mathcal{O}(k \cdot \rho \cdot \log n)$ in the LP rounding step. This represents our approximation ratio. The result may seem like a trivial generalization of the $k = 1$ case. However, this is not true as multiple graphs make the problem much harder. We can justify the approximation factor by a hardness bound. It shows that for asymmetric channels our algorithms are close to optimal without making further assumptions about the interference model.

THEOREM 5.1. *For each ρ, k with $\rho \cdot k = \mathcal{O}(\log n)$ there is no $\rho \cdot k / 2^{\mathcal{O}(\sqrt{\log(\rho \cdot k)})}$ approximation algorithm for asymmetric channels unless* P= NP.

PROOF. Again, such an algorithm could be used to approximate the independent set problem in bounded-degree graphs. Given a graph $G = (V, E)$ with maximum degree d, we construct k graphs $G_1 = (V, E_1)$, ..., $G_k = (V, E_k)$ each having an inductive independence number of $\rho = d/k$. For simplicity of notation, we assume this is an integer.

Let $\{v_1, \ldots, v_n\}$ be an arbitrary ordering of the vertices. We now distribute the edges from E to the edge sets E_1, ..., E_k. For a vertex v_i the incident edges to vertices v_j of lower index are distributed such that each edge set gets at most ρ such edges. Since the maximum vertex degree is d this is always possible. The valuations for the vertices are chosen such that for all vertices $b_{v,T}$ is 1 only for $T = [k]$ and 0 otherwise.

By this construction allocations of valuation b exactly correspond to independent sets of size b. Thus, such an approximation algorithm cannot exist unless P= NP. □

6. RECENT PROGRESS AND OPEN PROBLEMS

In this paper we present a general framework for secondary spectrum auctions addressing spectrum redistribution problems that represent a major obstacle to the development of wireless networking technology today. Our approach of using conflict graphs with bounded inductive independence number offers application to a wide variety of interference models. This allows to extend combinatorial auctions by assigning items/channels to *independent sets* of bidders.

Recently, we have improved the results for practically relevant classes of valuation functions $b(v, T)$ [Hoefer and Kesselheim 2012]. In these classes we can significantly reduce the dependence on the number k of channels. While for symmetric valuations we obtain an algorithm with a bound of $O(\rho \cdot (\log n + \log k))$, for submodular matroid-rank-sum valuations we even reach $O(\rho \cdot \log n)$, thereby dropping the dependence on k completely.

For general valuations, by spending a logarithmic factor we can strengthen the truthfulness guarantee to a notion based on stochastic dominance, which is stronger than truthfulness in expectation used in this paper and [Hoefer and Kesselheim 2012]. The algorithm presented in [Hoefer et al. 2013] uses a fractional oversampling technique that replaces the LP-optimum by a more structured solution, which is then decomposed and rounded by randomized meta-rounding.

The meta-rounding framework allows an elegant transformation of approximation algorithms into mechanisms that are truthful in expectation. However, it requires repeated application of the ellipsoid method, which is prohibitive in practice. Very recently, we have applied random sampling techniques [Hoefer and Kesselheim 2013] to obtain simple and quick mechanisms for the case of single-parameter valuations. While we lose a logarithmic factor in the approximation guarantee, the resulting mechanisms are even universally truthful, i.e., they remain truthful even if we publish the internal random coin flips of the algorithm in advance.

Some open problems in this domain are as follows. How can we extend known algorithms for (important classes of) ordinary combinatorial auctions or develop new techniques for secondary spectrum auctions? Our recent work shows promising first steps for techniques based on fractional oversampling and random sampling. Can we use other techniques for spectrum auctions, improve the guarantees, or obtain tight inapproximability results?

For truthfulness we use randomized meta-rounding and the ellipsoid method. Can this be avoided to make algorithms more applicable in practice? In particular, can we derive universally or deterministically truthful mechanisms with non-trivial guarantees that go beyond the special cases of [Hoefer and Kesselheim 2012; Hoefer and Kesselheim 2013]?

Another direction is to show that other interference models fit into our framework and prove tight bounds on the inductive independence number resulting from such models.

Finally, the inductive independence number itself represents a natural property of unweighted and weighted graphs that is still quite poorly understood.

REFERENCES

Akcoglu, K., Aspnes, J., DasGupta, B., and Kao, M.-Y. 2002. Opportunity cost algorithms for combinatorial auctions. In *Applied Optimization 74: Computational Methods in Decision-Making, Economics and Finance*, E. Kontoghiorghes, B. Rustem, and S. Siokos, Eds. Kluwer, 455–479.

Alicherry, M., Bhatia, R., and Li, L. E. 2006. Joint channel assignment and routing for throughput optimization in multi-radio wireless mesh networks. *IEEE J. Sel. Areas Commun. 24,* 11, 1960–1971.

Balakrishnan, H., Barrett, C., Kumar, A., Marathe, M., and Thite, S. 2004. The distance-2 matching problem and its relationship to the MAC-layer capacity of ad hoc wireless networks. *IEEE J. Sel. Areas Commun. 22,* 6, 1069–1079.

Barrett, C., Kumar, A., Marathe, M., Thite, S., and Istrate, G. 2006. Strong edge coloring for channel assignment in wireless radio networks. In *Proc. 4th Pervasive Computing and Communications Workshops (PERCOMW).* 106.

Buddhikot, M., Kolodzy, P., Miller, S., Ryan, K., and Evans, J. 2005. DIMSUMNet: New directions in wireless networking using coordinated dynamic spectrum access. In *Proc. IEEE WoWMoM.* 78–85.

Christodoulou, G., Elbassioni, K., and Fouz, M. 2010. Truthful mechanisms for exhibitions. In *Proc. 6th Intl. Workshop Internet & Network Economics (WINE).* 170–181.

Cramton, P., Shoham, Y., and Steinberg, R., Eds. 2006. *Combinatorial Auctions.* MIT Press.

Dobzinski, S., Nisan, N., and Schapira, M. 2012. Truthful randomized mechanisms for combinatorial auctions. *J. Comput. Syst. Sci. 78,* 1, 15–25.

Fanghänel, A., Geulen, S., Hoefer, M., and Vöcking, B. 2013. Online capacity maximization in wireless networks. *J. Sched. 16,* 1, 81–91.

Fanghänel, A., Kesselheim, T., Räcke, H., and Vöcking, B. 2009. Oblivious interference scheduling. In *Proc. 28th Symp. Principles of Distrib. Comput. (PODC).* 220–229.

Fanghänel, A., Kesselheim, T., and Vöcking, B. 2011. Improved algorithms for latency minimization in wireless networks. *Theoret. Comput. Sci. 412,* 24, 2657–2667.

Feige, U. and Vondrák, J. 2006. Approximation algorithms for allocation problems: Improving the factor of 1-1/e. In *Proc. 47th Symp. Foundations of Computer Science (FOCS).* 667–676.

Fishkin, A. V. 2003. Disk Graphs: A Short Survey. In *Proc. 1st Intl. Workshop Approx. and Online Algorithms (WAOA).* 260–264.

Gandhi, S., Buragohain, C., Cao, L., Zheng, H., and Suri, S. 2008. Towards real-time dynamic spectrum auctions. *Computer Networks 52,* 4, 879–897.

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.

GRÄF, A., STUMPF, M., AND WEISSENFELS, G. 1998. On coloring unit disk graphs. *Algorithmica 20,* 3, 277–293.

GUPTA, P. AND KUMAR, P. R. 2000. The capacity of wireless networks. *IEEE Trans. Inf. Theory 46,* 2, 388–404.

HALLDÓRSSON, M. 2013. Wireless scheduling with power control. *ACM Trans. Algorithms 9,* 1, 7.

HALLDÓRSSON, M., HOLZER, S., MITRA, P., AND WATTENHOFER, R. 2013. The power of non-uniform wireless power. In *Proc. 24th Symp. Discrete Algorithms (SODA)*. 1595–1606.

HALLDÓRSSON, M. AND MITRA, P. 2011. Wireless capacity with oblivious power in general metrics. In *Proc. 22nd Symp. Discrete Algorithms (SODA)*. 1538–1548.

HÅSTAD, J. 1999. Clique is hard to approximate within $n^{1-\varepsilon}$. *Acta Mathematica 182,* 1, 105–142.

HOCHBAUM, D. 1983. Efficient bounds for the stable set, vertex cover and set packing problems. *Disc. Appl. Math. 6,* 3, 243–254.

HOEFER, M. AND KESSELHEIM, T. 2012. Secondary spectrum auctions for symmetric and submodular bidders. In *Proc. 13th Conf. Electronic Commerce (EC)*. 657–671.

HOEFER, M. AND KESSELHEIM, T. 2013. Brief announcement: Universally truthful secondary spectrum auctions. In *Proc. 25th Symp. Parallelism in Algorithms and Architectures (SPAA)*. 99–101.

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.

HOEFER, M., KESSELHEIM, T., AND VÖCKING, B. 2013. Truthfulness and stochastic dominance with monetary transfers. In *Proc. 14th Conf. Electronic Commerce (EC)*. 567–582.

HOLZMAN, R., KFIR-DAHAV, N., MONDERER, D., AND TENNENHOLTZ, M. 2004. Bundling equilibrium in combinatorial auctions. *Games Econom. Behav. 47,* 1, 104–123.

ILERI, O., SAMARDZIJA, D., SIZER, T., AND MANDAYAM, N. 2005. Demand responsive pricing and competitive spectrum allocation via a spectrum server. In *Proc. 1st Symp. Dynamic Spectrum Access Networks (DySPAN)*. 194–202.

KAKO, A., ONO, T., HIRATA, T., AND HALLDÓRSSON, M. 2009. Approximation algorithms for the weighted independent set problem in sparse graphs. *Disc. Appl. Math. 157,* 4, 617–626.

KASH, I., MURTY, R., AND PARKES, D. 2011. Enabling spectrum sharing in secondary market auctions. In *6th Workshop Econom. Networks, Syst. and Comput. (NetEcon)*.

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. 2012. Approximation algorithms for wireless link scheduling with flexible data rates. In *Proc. 20th European Symp. Algorithms (ESA)*. 659–670.

KESSELHEIM, T. AND VÖCKING, B. 2010. Distributed contention resolution in wireless networks. In *Proc. 24th Intl. Symp. Distrib. Comput. (DISC)*. 163–178.

KRUMKE, S., MARATHE, M., AND RAVI, S. S. 2001. Models and approximation algorithms for channel assignment in radio networks. *Wireless Networks 7,* 6, 575–584.

LAVI, R. AND SWAMY, C. 2011. Truthful and near-optimal mechanism design via linear programming. *J. ACM 58,* 6, 25.

NISAN, N., TARDOS, É., ROUGHGARDEN, T., AND VAZIRANI, V., Eds. 2007. *Algorithmic Game Theory.* Cambridge University Press.

TREVISAN, L. 2001. Non-approximability results for optimization problems on bounded degree instances. In *Proc. 33rd Symp. Theory of Computing (STOC)*. 453–461.

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.

YE, Y. AND BORODIN, A. 2012. Elimination graphs. *ACM Trans. Algorithms 8,* 2, 14.

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.

ZHU, Y., LI, B., AND LI, Z. 2012. Truthful spectrum auction design for secondary networks. In *Proc. 31st IEEE Conf. Computer Communications (INFOCOM)*. 873–881.

A. MECHANISM DESIGN

In this section we show how to apply the randomized meta-rounding framework [Lavi and Swamy 2011] to obtain a truthful mechanism for the problem, in which the valuations for the allocations are private information. We only highlight the main ideas of this technique and the most important observations that allow the use for our problem.

The main idea of the approach is to decompose an optimal LP solution x^* into a set of polynomially many integral solutions with the following property. For each integral solution we determine a probability, and the expected social welfare of a randomly chosen solution according to the probabilities is exactly b^*/α, where in our case $\alpha = 8 \cdot \sqrt{k} \cdot \rho$. Given such a decomposition, we can use scaled VCG payments to implement a randomized mechanism that is truthful in expectation. For an accessible presentation of the complete technique, see [Nisan et al. 2007, Chapter 12] or [Lavi and Swamy 2011].

In particular, for simplicity let us first consider only a constant number of channels; the adjustment to arbitrary many channels is treated below. We ask the vertices to obtain all valuations for all channel bundles and solve the corresponding LP (interference information is assumed to be publicly available). Note that at this point we are given the optimal solution to an *infeasible* LP. We set up a decomposition LP with exponentially many variables – one for each *feasible* integral solution – that represent our desired probabilities. This LP has polynomially many constraints but exponentially many variables. We can construct the dual with polynomially many variables and exponentially many constraints. The variables can be interpreted as valuations in an adjusted combinatorial auction problem. If this problem has an algorithm that verifies an integrality gap, we obtain a separation oracle and can solve the dual decomposition LP in polynomial time. In particular, it allows us to construct an equivalent LP with a polynomial number of constraints, i. e., the ones corresponding to the solutions obtained by our algorithm. For this polynomial-sized dual we construct the primal and determine the polynomially many probabilities of the solutions found by our algorithm, which completes the decomposition.

It remains to verify that our algorithms provide integral solutions within the desired integrality gap of α for the adjusted combinatorial auction problems using dual variables as valuations. We note here that our algorithms bound the integrality gap of LP (1) and (4), and they can be derandomized using the technique of pairwise independence. In this way, given an optimal LP solution x^* we can obtain an integral solution of value at least b^*/α. Note that our LP describe, in fact, relaxations of the combinatorial auction problem with conflict graphs, because Conditions (1b) and (4b) allow each vertex to have multiple neighbors on the same channel. An arbitrary integral solution to the LP might thus be infeasible for the original problem. This is even more severe in the case of the physical model with power control, where even the interpretation of edge weights is significantly disconnected from the actual interference that is received. However, our algorithms produce feasible integral solutions with the desired gap to the infeasible fractional optimum. Thus, they also prove the gap for a potential fractional optimum to the LP describing the (more constrained) exact combinatorial auction problem with conflict graphs in the respective cases. The remaining arguments can be adapted from [Lavi and Swamy 2011] almost without adjustment.

In case of an arbitrary number of channels, we can use demand oracles to solve the LPs. This results in only a polynomial number of (non-zero) variables for the LP and for the dual of the decomposition LP. Note that the procedure to separate the dual of the decomposition LP does not require demand oracles. In fact, the complete decomposition procedure can be carried out without accessing the original bidder valuations.

B. BOUNDING THE INDUCTIVE INDEPENDENCE NUMBER IN INTERFERENCE MODELS

Another example for the transmitter scenario is the so-called *distance-2 coloring*. In contrast to the above model not only the neighbors (with intersecting disks) must be on different channels but also their neighbors. Distance-2 coloring is a common model of transmitter scenarios. Here, we analyze the restriction on two graph classes. We refer the reader to [Krumke et al. 2001] for the exact definitions and a discussion of the model. We can prove $\rho = \mathcal{O}(1)$ as well in this case.

LEMMA B.1. *Let $r > 0$, $a > 0$ and D be a disk of radius ar. Then the number of disks of radius at least r that intersect D but not each other is at most $(a + 2)^2$.*

PROOF. W. l. o. g., we assume the surrounding disks to have radius exact r. By scaling them down and moving them inside their original area, they still do not intersect each other. By moving them to the respective closest location to D, they still intersect D.

The disks of radius r are fully contained within the disk of radius $kr + 2r$ around the center of D. Each takes an area of πr^2, whereas the available area is only $\pi(ar + 2r)^2$. So, the number of surrounding disks is at most $\pi(ar + 2r)^2/\pi r^2 = (a + 2)^2$. □

PROPOSITION B.2. *For Distance-2 coloring in disk graphs the associated conflict graph has an inductive independence number $\rho = \mathcal{O}(1)$.*

PROOF. As for disk graphs, we order the vertices by decreasing ranges. Now consider a vertex v and a conflicting vertex u of larger range. This vertex can either be directly connected to v (there are at most 5 ones of this kind) or via an intermediate vertex u'. If the $r_{u'} < r_v$ is smaller, we see that the disk of radius r_u around u intersects the one of radius $2r_v$ around v. The above lemma yields that there can be at most a constant number of such vertices. For the case $r_{u'} \geq r_v$, we take into consideration the disks around the intermediate vertices also do not intersect. So, there can be at most 5 intermediate vertices and as many conflicting vertices. The total number of conflicting vertices is constant. □

PROPOSITION B.3. *For Distance-2 coloring in (r, s)-civilized graphs the inductive independence number of the associated conflict graph is $\rho \leq (4r/s + 2)$.*

PROOF. In this case, the ordering does not matter. Therefore, we do not need to know the geometric representation of the graph.

Consider a vertex v and a set of vertices M conflicting with v but not with each other. Since the path length from v to each vertex in M is at most 2, the distance in the plane is at most $2r$. Now consider disks around the vertices in M, each of radius $s/2$. By definition of the (r, s)-civilized graph these disks do not intersect each other. However, each of them intersects a disk of radius $2r$ around v. Applying the above lemma, we see there are at most $(4r/s + 2)^2$ such disks. □

As a matter of fact ρ has to depend on this ratio of r and s. Obviously, all graphs can be represented as (r, s)-civilized if the ratio r and s is unbounded. However, our algorithm's running time does not depend on r and s. Therefore, the approximation factor has to depend on them.

Finally, we discuss another link-based approach of *distance-2 matching* [Balakrishnan et al. 2004]. In this case, two edges $e \neq e'$ may be allocated to the same channel if there are at least two edges on any connecting path. Typically, results are restricted to certain graph classes, because in general approximating maximum distance-2 matchings is hard. For disk graphs, we can also show that the corresponding conflict graph has $\rho = \mathcal{O}(1)$. Interestingly, for distance-2 matching there is already an algorithm and analysis based on the observation that the inductive independence number is bounded, but the concepts are termed differently. Barrett et al. [Barrett et al. 2006] analyze a greedy approach to find a maximum independent set. For a link $e = (u, v)$, they define $r(e) = r(u) + r(v)$, where $r(u)$ and $r(v)$ are the radius of the disk surrounding u resp. v. The algorithm orders the links by increasing values of $r(e)$. The key observation is now that for all links e the maximum number of links of higher index that collide with e but not with each other is $\mathcal{O}(1)$. This immediately yields $\rho = \mathcal{O}(1)$.

COROLLARY B.4. *For distance-2 matching in disk graphs the associated conflict graph has an inductive independence number* $\rho = \mathcal{O}(1)$.

Analyses of greedy algorithms are often carried out in a similar manner. Such arguments already suffice to bound the inductive independence number. There is plenty of opportunity to further extend our results by similar observations.

C. THEOREMS USED IN THE PROOF OF THEOREM 4.5

Theorems 1 and 7 in [Kesselheim 2011] provide necessary conditions for *admissible sets*. These are sets of links, for which there is a power assignment, such that the SINR constraint is fulfilled.

THEOREM C.1 (THEOREM 1 FROM [KESSELHEIM 2011]). *Let* $\mathcal{L}$ *be an admissible set. Let furthermore* s *and* r *be some arbitrary nodes in* V. *Then we have*

$$\sum_{\substack{(s',r')\in\mathcal{L}\\ d(s',r')\geq d(s,r)}} \min\left\{1, \frac{d(s,r)^\alpha}{d(s,r')^\alpha}\right\} + \min\left\{1, \frac{d(s,r)^\alpha}{d(s',r)^\alpha}\right\} = \mathcal{O}(1)$$

in fading metrics[2].

THEOREM C.2 (THEOREM 7 FROM [KESSELHEIM 2011]). *Let* $\mathcal{L}$ *be an admissible set. Let furthermore* s *and* r *be some arbitrary nodes in* V. *Then we have*

$$\sum_{\substack{(s',r')\in\mathcal{L}\\ d(s',r')\geq d(s,r)}} \min\left\{1, \frac{d(s,r)^\alpha}{d(s,r')^\alpha}\right\} + \min\left\{1, \frac{d(s,r)^\alpha}{d(s',r)^\alpha}\right\} = \mathcal{O}(\log|\mathcal{L}|)$$

in any metric.

Theorem 3 in [Kesselheim 2011] gives a sufficient condition for a set to be admissible. In more detail, it states a condition under which the power-control procedure stated in [Kesselheim 2011] succeeds. Instead of using this procedure, the power assignment could also be computed by solving the linear equation system directly.

THEOREM C.3 (THEOREM 3 FROM [KESSELHEIM 2011]). *If in a link set* $\mathcal{L}$ *the following condition is fulfilled for all* $\ell' \in \mathcal{L}$

$$\sum_{\substack{(s,r)\in\mathcal{L}\\ d(s,r)<d(s',r')}} \frac{d(s,r)^\alpha}{d(s,r')^\alpha} + \frac{d(s,r)^\alpha}{d(s',r)^\alpha} \leq \tau \qquad \textit{where } \tau = \frac{1}{2\cdot 3^\alpha \cdot (4\beta+2)}\ , \tag{7}$$

then the power-control procedure in [Kesselheim 2011] computes a power assignment fulfilling the SINR condition.

[2]*Fading metrics* have been introduced by Halldórsson [2013] and are a generalization of the standard assumption to consider the Euclidean plane im combination with $\alpha > 2$.