

# Universally Truthful Secondary Spectrum Auctions

Martin Hoefer\*

Thomas Kesselheim†

## Abstract

We present algorithms for implementing local spectrum redistribution in wireless networks using a mechanism design approach. For example, in single-hop request scheduling, secondary users are modeled as rational agents that have private utility when getting assigned a channel for successful transmission. We present a rather simple algorithmic technique that allows to turn existing and future approximation algorithms and heuristics into truthful mechanisms for a large variety of networking problems. In contrast to previous work, our approach works for virtually all known interference models in the literature, including the physical model of interference based on SINR. It allows to address single-hop and multi-hop scheduling, routing, and even more general assignment and allocation problems. Our mechanisms are randomized and represent the first universally-truthful mechanisms for these problems with rigorous worst-case guarantees on the solution quality. In this way, our mechanisms can be used to obtain guaranteed solution quality even with risk-averse or risk-seeking bidders, for which existing approaches fail.

---

\*Max-Planck-Institut für Informatik and Saarland University, Saarbrücken, Germany, [mhoefer@cs.rwth-aachen.de](mailto:mhoefer@cs.rwth-aachen.de). Supported by DFG through Cluster of Excellence MMCI.

†Dept. of Computer Science, Cornell University, Ithaca, NY, USA. [kesselheim@cs.cornell.edu](mailto:kesselheim@cs.cornell.edu). Supported by a fellowship within the Postdoc-Programme of the German Academic Exchange Service (DAAD) and by DFG through UMIC Research Center at RWTH Aachen University.

# 1 Introduction

A major challenge in the development of future wireless networking technology lies in spectrum scarcity, i.e., the limited availability of frequency bands for existing and new devices or services. To a large extent this problem results from the static allocation and licensing rules currently in use, where licenses for frequency bands are given to large service providers for entire countries. There is a major research effort underway in computer science and engineering to overcome this static allocation approach. The main idea is to motivate the licensed *primary users* to open up their bands in local areas and enable *secondary users* to use these local spectrum opportunities.

An attractive approach to implement secondary usage are *auctions*. There is a rich theory in economics showing that auctions allow to implement global objectives in a market with rational participants and private information. Auction mechanisms allow to implement secondary spectrum usage as a market, in which primary users can sell access to their unused licensed spectrum bands to secondary users [3]. In this scenario, secondary users are assumed to be selfish and rational. They can manipulate the outcome of a mechanism by misreporting their valuation and try to obtain a desired channel using smaller payments. Therefore, we desire a *truthful* mechanism that computes allocation and payments in a way that no user has an incentive to lie. It ensures that every user maximizes his utility by reporting the true valuation. For decades, the most prominent technique to design truthful mechanisms was VCG [5, 11, 23], which applies only if the chosen allocation optimizes social welfare exactly. Unfortunately, for non-trivial spectrum auctions social welfare maximization is NP-hard. The challenge is thus to find mechanisms that (1) are truthful, (2) run in polynomial time, and (3) provide allocations with near-optimal social welfare.

**Our Contribution** We introduce a unified framework to design simple polynomial-time mechanisms for a large variety of problems within secondary spectrum usage and prove non-trivial worst-case guarantees on their social welfare. We heavily extend the current literature on secondary spectrum auctions in several ways.

First, many existing secondary spectrum auctions model interference as a binary property using, e.g., unit-disk graphs, interference boundaries, or are otherwise based on local binary conflicts [9, 25–27]. These approaches lack many important features of realistic signal propagation. Instead, we here use realistic interference models based on the signal-to-interference-plus-noise ratio (SINR), for which allocation problems are often significantly more complicated to analyze.

Second, while there exist approaches to secondary spectrum auctions with SINR [15, 16], they are mostly unsuitable for practical use. Being polynomial-time algorithms in principle, their main drawback is the time-intensive use of randomized meta-rounding [4, 19] and the ellipsoid method for convex optimization. The meta-rounding technique relies on the ellipsoid method and iterative application of separation oracles. This approach is used by virtually all known secondary spectrum auctions with non-trivial performance guarantees, even the ones for binary interference models [9, 27]. In contrast, our mechanisms are fast, surprisingly simple to grasp and implement, without complicated convex optimization techniques. Nevertheless, we prove non-trivial worst-case bounds on their allocation quality.

Third, the meta-rounding approach used in related work yields randomized mechanisms that are *truthful in expectation*, i.e., if users care only about the *expectation* of utility, they have no incentive to lie. Truthfulness in expectation is a strong assumption, because users must be risk-neutral. Such mechanisms lose their truthfulness if users are, e.g., risk-averse or risk-seeking or take further parameters of the utility distribution into account. In contrast, our mechanisms are

*universally truthful*, randomization has no effect on incentives. Our mechanisms can be thought of as probability distributions over deterministic truthful mechanisms. We use randomization only to guarantee good social welfare.

Fourth, while we present our approach for secondary markets with single-hop transmissions in the SINR model, it turns out to be much more flexible. In fact, it is applicable to a large variety of mechanism design problems in (wireless) networks. The main criterion is that feasibility in the underlying allocation problem is downward closed. This holds, e.g., for all allocation problems that can be modeled via packing linear or integer programs. We highlight this extension of our approach for multi-hop routing problems with binary [27] or even SINR interference.

The general idea of our approach is as follows. Similar to [6], we first determine a “representative” valuation of the users by independently at random moving each user into a statistics group with small probability. The users in this group are asked their valuation and discarded without channel and payments. Based on their reports, a representative valuation is determined, from which we obtain a random take-it-or-leave-it price  $p$ . Each surviving user is asked if he would want to buy a channel at price  $p$ . If not, the user leaves the auction, otherwise he survives. For the set of surviving users, we then determine an allocation using an arbitrary polynomial-time approximation algorithm for the channel allocation problem. Finally, each user that ends up with a channel pays the price  $p$ , all others pay nothing. While truthfulness of this scheme can be shown rather directly, the challenge is proving that this general framework allows to obtain non-trivial worst-case guarantees on social welfare and revenue.

**Related Work** In recent years, there have been numerous studies of different flavors of the problem of maximizing the number of successful transmissions in a wireless network with interference. Early works model this problem essentially as a maximum independent set problem in graphs [7, 22]. However, as physical conditions are only poorly captured this way, starting with a paper by Moscibroda and Wattenhofer [20] attention has shifted to models based on SINR. Many approximation algorithms use fixed power assignments, e.g., uniform ones [2, 14], or, more generally, monotone ones [12, 13]. Better results in terms of the approximation factor can, however, be achieved by making also the power assignment itself subject to optimization [17, 18]. To illustrate our techniques, we use the latter algorithm, which is the best algorithm to solve the combined selection and power-control problem so far. However, we also show that our approach is much more general, allowing it to be used with any of the algorithms mentioned above as well as any future algorithm.

Designing auction mechanisms for secondary spectrum markets has attracted increased research interest in recent years. Usually, socially optimal channel assignment poses NP-hard graph coloring and maximum independent set problems. In these cases, the classic approach of VCG for designing truthful mechanisms [5, 11, 23] cannot be applied efficiently. Most approaches to spectrum auctions are single-parameter problems, where each user has a single numerical value for getting any one of a set of channels. They are often based on the classic monotonicity characterization of truthful mechanisms [21]. Zhou et al [25] studied monotone deterministic algorithms for a secondary spectrum auction in a graph-based interference model. This was extended by Zhou and Zheng [26] to a double auction framework. More recently, Gopinathan et al [9] studied truthful-in-expectation mechanisms via randomized meta-rounding that allow to bound the worst-case approximation factor and allow to incorporate additional objectives like fairness. In a related work, Gopinathan et al [8] also studied revenue optimization using concepts introduced for revenue-optimizing auctions of digital goods. Very recently, Zhu et al [27] proposed deterministic monotone mechanisms for a single-parameter multi-hop routing model. They also provided truthful-in-expectation mechanisms

based on randomized meta-rounding with provable worst-case performance.

All the above mentioned works address single-parameter domains. The general case, in which each user can have a different valuation for each subset of acquired channels, was studied in our work [16] for single-hop scheduling in a general framework encompassing graph-based and SINR interference models. Very recently, we provided improved results in the same framework for the popular subclasses of symmetric and submodular user valuations [15]. In both papers, we proposed truthful-in-expectation mechanisms based on randomized meta-rounding with provable worst-case guarantees. Our guarantees depend on the inductive independence number [1], a non-standard graph parameter that can capture feasibility in virtually all known interference models.

**Outline** The remainder of this paper is structured as follows. We define the initial scenario and mechanism design basics in Section 2. In Section 3 we describe our mechanism for single-hop request scheduling in the SINR model. In Section 4, we then show how to generalize our technique to many other problems, including single- and multi-hop scheduling in the SINR and other interference models. Finally, we conclude in Section 6.

## 2 Preliminaries

**Network Model** We will first explain our approach in the context of single-hop wireless transmission scheduling and then generalize our technique later on. We assume there is a primary user that holds a set  $\mathcal{C}$  of  $k$  channels in a local area. There is a set  $\mathcal{N}$  of  $n$  secondary users that strive to obtain a channel. Each user corresponds to a communication request or *link* between points in a metric space. Link  $i$  consists of sender  $s_i$  and receiver  $r_i$ .

For each channel, the primary user assigns licenses to allow a subset  $\mathcal{L} \subseteq \mathcal{N}$  of users to use the channel. If link  $i$  gets a license, sender  $s_i$  transmits on the channel a signal with power  $\sigma_i$ . We will consider different power assignments below. The signal strength decays exponentially over distance with exponent  $\alpha > 1$ . Given the subset  $\mathcal{L} \subseteq \mathcal{N}$  of links and a power assignment  $\sigma : \mathcal{L} \rightarrow \mathbb{R}_{>0}$ , the transmission on link  $i \in \mathcal{L}$  is successful if the signal-to-interference-plus-noise ratio (SINR)  $\gamma_i(\mathcal{L}, p)$  is above some predefined threshold  $\beta > 0$ :

$$\gamma_i(\mathcal{L}, \sigma) = \frac{\sigma_i / d_{ii}^\alpha}{N + \sum_{j \in \mathcal{L}, j \neq i} \sigma_j / d_{ji}^\alpha} \geq \beta, \quad (1)$$

where  $d_{ji}$  is the distance between sender  $s_j$  and receiver  $r_i$  and  $N$  is ambient noise. This condition captures the intuition that at  $r_i$  the decayed signal of  $s_i$  is significantly louder than interference from decayed signals of other senders and noise.

Each user  $i$  has a benefit  $w_i$  for being able to transmit successfully on a channel. In particular, its benefit is  $w_i > 0$  if and only if he gets a license for at least one channel on which condition (1) is fulfilled. Otherwise, its benefit is 0. The goal of the primary user is to compute an allocation  $\mathcal{A} = (\mathcal{L}_1, \dots, \mathcal{L}_k)$  of secondary users in  $\mathcal{N}$  to channels to maximize *social welfare*, i.e., the sum of benefits of successful users.

**Mechanism Design** Benefits are private information of secondary users. Before the allocation each secondary user  $i$  must report its benefit  $w_i$  for a successful transmission; we say user  $i$  makes *bid*  $b_i$ . Users are rational and selfish, they would like to obtain a channel at the lowest possible

cost. Depending on the channel allocation algorithm, user  $i$  might benefit from misreporting its value and have an incentive to lie  $b_i \neq w_i$ .

This fundamental problem has been studied for several decades in the area of *mechanism design*. To avoid manipulation and set incentives for truthful bids, we design a *truthful mechanism*. It collects the vector of bids  $b = (b_i)_{i \in \mathcal{N}}$ , allocates users to channels, and charges payments  $p_i(b)$  to user  $i \in \mathcal{N}$ . User  $i \in \mathcal{N}$  has a quasi-linear *utility*

$$u_i(b) = \begin{cases} w_i - p_i(b) & (1) \text{ holds on at least one channel assigned to } i \\ 0 & \text{otherwise} \end{cases}.$$

A deterministic mechanism is *truthful* if no user profits from lying. More formally, we require that

$$u_i(b_i, b_{-i}) \leq u_i(w_i, b_{-i}) \quad \text{for all } b_i \geq 0 \text{ and } b_{-i}, \quad (2)$$

where we use  $b_{-i}$  to denote the vector  $b$  of bids excluding bid  $b_i$ . The classic technique to compose deterministic truthful mechanisms is VCG [5, 11, 23], but it applies only when the allocation maximizes social welfare exactly. Maximizing social welfare in the channel allocation problem is NP-hard even for special cases [2]. We have to find other ways to obtain truthful mechanisms with good social welfare guarantees that run in polynomial time.

Towards this end, we study randomized mechanisms, which are *truthful in expectation* if (2) holds for the *expected* utility of every user. Truthfulness in expectation has the drawback that the outcome of random coins must be invisible to the users, and except for the expectation every user must be indifferent to all other parameters of its utility distribution.

A much stronger condition is fulfilled if a randomized mechanism is *universally truthful*. We can interpret a universally truthful mechanism as having a number of deterministic truthful mechanisms and initially making one randomized decision which one of the deterministic mechanisms to apply. Here truthfulness is independent of randomization, so the mechanism could make the random decision in public before collecting bids. The reason we use randomization is to obtain a better approximation of social welfare.

### 3 Universally Truthful Auctions for Single-Hop SINR Scheduling

The general structure of our mechanism is presented in Algorithm 1. It is inspired by a universally truthful mechanism for combinatorial auctions by Dobzinski et al. [6]. At the beginning, the mechanism decides at random whether the actual allocation algorithm is executed at all. The purpose of this step is to address the case when there is a super-dominant bidder whose bid is much higher than the sum of valuations of all other bidders. To treat this case, the mechanism runs a second-price auction in which only a single bidder is served with a small, constant probability  $\varepsilon$ . In this case, only the bidder having the highest bid is allocated a channel and is charged the second highest bid.

With probability  $1 - \varepsilon$ , the actual allocation algorithm is performed. We consider prior-free mechanisms without any knowledge on the bidders' valuations. Therefore, we sample a number of bidders that are asked for their valuation and discarded afterwards. If the auction mechanism is applied repeatedly in practice, knowledge about typical bidder valuations might become available due to experience, which allows to skip this sampling step. Having obtained knowledge on "typical" valuations, we determine a price in a random experiment. The price  $p$  is obtained by choosing  $X$  uniformly at random from  $\{0, 1, \dots, \lceil \log n \rceil + 1\}$  and setting  $p := 2^{-X} \cdot B$ , where  $B$  is the largest bid

---

**Algorithm 1:** Universally-truthful framework

---

```
1 set SEC-PRICE to 1 with probability  $\varepsilon$ , otherwise to 0;
2 if SEC-PRICE = 1 then
3   Let  $i^*$  be a bidder such that  $b_{i^*} = \max_{i \in \mathcal{N}} b_i$ ;
4   Pick an arbitrary channel in  $\mathcal{C}$  and allocate it to  $i^*$ ;
5   Charge  $p_{i^*}(b) := \max_{i \in \mathcal{N} \setminus \{i^*\}} b_i$ ;
6   All other bidders  $i$  are not allocated and charged  $p_i(b) := 0$ ;
7 else
8   forall the  $i \in \mathcal{N}$  do
9      $\perp$  add  $i$  to STAT with probability  $\varepsilon$ , otherwise to FIXED
10  Set  $B := \max_{i \in \text{STAT}} b_i$ ;
11  choose  $X$  uniformly at random from  $\{0, 1, \dots, \lceil \log n \rceil + 1\}$ ;
12  Set  $p := 2^{-X} \cdot B$ ;
13  Set  $M := \{i \in \text{FIXED} \mid b_i \geq p\}$ ;
14  Run UnweightedPacking( $M, \mathcal{C}$ ) which returns allocation  $(\mathcal{L}_1, \dots, \mathcal{L}_k)$ ;
15  If  $i \in \bigcup_{j=1}^k \mathcal{L}_j$ , charge  $p_i(b) := p$ , otherwise  $p_i(b) := 0$ ;
```

---

that was observed in the sampling phase. Afterwards, we only consider bidders that are willing to pay at least  $p$ . For these bidders, we compute an allocation, in which each bidder is either selected and assigned to a channel in  $\mathcal{C}$  or it is discarded (see below). In the end, each bidder that is selected and assigned to a channel is charged  $p$ . The discarded bidders are not charged.

The selection of bidders and allocation to channels is delegated to a procedure called **UnweightedPacking**( $M, \mathcal{C}$ ). Given a set  $M \subseteq \mathcal{N}$  of bidders, it computes a partial allocation of bidders in  $M$  to channels in  $\mathcal{C}$  that is feasible with respect to the SINR constraints. Note that the routine does not get to know the valuations of bidders, instead it aims at maximizing only the *number* of selected bidders. Their actual valuations are neglected at this stage. We here assume that the application scenario allows power control, which allows us to employ Algorithm 2 originally presented in [17]. It is known that this algorithm always computes feasible allocations. In particular, the desired power assignment fulfilling the SINR constraints of all selected links in Line 5 exists. Furthermore, it is a constant-factor approximation algorithm for the optimization problem of maximizing the number of feasible requests on  $k$  channels (for a detailed treatment see [17]). In our case, this implies that for a given set  $M$ , we have the following property: Let  $S = \bigcup_{j=1}^k \mathcal{L}_j$  be the set of bidders selected on any channel after executing **UnweightedPacking**( $M, \mathcal{C}$ ). Let furthermore  $F \subseteq M$  be any arbitrary set of bidders such that there is a feasible allocation of all bidders in  $F$  to the channels in  $\mathcal{C}$ . Then we have  $|S| \geq \kappa \cdot |F|$  for some constant  $\kappa \leq 1$ .

If the application scenario does not allow power control, we can alternatively implement **UnweightedPacking**( $M, \mathcal{C}$ ) by adapting algorithms for fixed power assignments. We discuss this issue as part of the general treatment in the next section.

We will first show that Algorithm 1 is a universally truthful mechanism. Given any fixed outcome of the random experiments, no bidder can profit from misreporting the valuation.

**Theorem 1.** *Algorithm 1 is universally truthful.*

*Proof.* We have to show that the mechanism is a distribution over (deterministic) truthful mechanisms. Let us now fix the outcome of *SEC-PRICE*. In case of *SEC-PRICE* = 0, fix furthermore

all random choices leading to the division into STAT and FIXED, and also the random choice of  $X$  leading to  $p$ .

Now, consider a single bidder  $i \in \mathcal{N}$ . We have to show that she cannot gain by misreporting the valuation. For this purpose, we distinguish between the outcomes of the randomization.

**Case 1:  $SEC\text{-}PRICE = 1$**  In this case, we perform the classic deterministic second-price Vickrey auction for a single item [23] that is known to be monotone [21] and truthful.

**Case 2:  $SEC\text{-}PRICE = 0$  and  $i \in \text{STAT}$**  Bidder  $i$  will not be allocated anything, independent of the bid. Thus, she has no incentive to misreport.

**Case 3:  $SEC\text{-}PRICE = 0$  and  $i \in \text{FIXED}$**  The bid  $b_i$  only determines whether or not  $i$  is contained in set  $M$ . The price  $p$  cannot be influenced by  $i$ , because it is determined only by the bidders in STAT. Now, let us assume that  $i$  reports a false bid  $b'_i$ , which results in a different allocation  $S'$ .

The two cases that either  $b_i, b'_i \geq p$  or  $b_i, b'_i \leq p$  are irrelevant as in both cases  $M$ , the allocation and the payment are the same for  $b_i$  and  $b'_i$ . Hence, it remains to consider the cases  $b_i < p \leq b'_i$  and  $b'_i < p \leq b_i$ . In the first case,  $i$  can be contained in  $S'$  but not in  $S$ . If  $i \notin S'$ , she has zero utility 0 for both bids. If  $i \in S'$ , then bidder  $i$  has to pay  $p$ , which is more than her valuation. Hence, her quasi-linear utility is negative as opposed to zero utility for bid  $b_i$ . In the case  $b'_i < p \leq b_i$ , we have in turn that  $i \in S$  but  $i \notin S'$ . In case of bidding  $b_i$ , the utility is  $b_i - p \geq 0$ , whereas in case of bidding  $b'_i$  it is 0. In any case, the utility is maximized by bidding  $b_i$ .  $\square$

This shows that our algorithm is a truthful mechanism independent of random choices. We now proceed to the proof that the allocation returned by the mechanism provides a non-trivial worst-case approximation guarantee. Our proof proceeds in two cases. If the bid of the highest bidder is very large compared to the optimal social welfare, then we resort to the single-item Vickrey auction executed with some probability that gives him a channel. In this way, our algorithm secures the high bid with significant probability. Otherwise, we bound the social welfare of our allocation using the total revenue generated through the prices  $p$  paid by the allocated bidders. As no bidder gets negative utility, the revenue is a lower bound for the welfare of the allocation obtained by the algorithm.

**Theorem 2.** *Let  $\mathcal{A}$  be the allocation returned by the mechanism and  $\mathcal{A}^*$  be the allocation optimizing social welfare. Then we have*

$$\mathbf{E}[b(\mathcal{A})] = \Omega\left(\frac{\varepsilon}{\log n}\right) \cdot b(\mathcal{A}^*) .$$

*Proof.* We denote the highest bid by  $B^* = \max_{i \in \mathcal{N}} b_i$ . Note that the channels are important for feasibility of the solution but not for the social welfare. In particular, a bidder is indifferent about

---

**Algorithm 2:** UnweightedPackingPC( $M, \mathcal{C}$ ) [17]

---

- 1 initialize  $\mathcal{L}'_1, \dots, \mathcal{L}'_k = \emptyset$  ;
  - 2 **for**  $\ell' \in \mathcal{R}$  *in order of increasing length* **do**
  - 3     **if** *there is some  $j \in \mathcal{C}$  such that  $\sum_{\substack{(s,r) \in \mathcal{L}'_j \\ d(s,r) < d(s',r')}} \frac{d(s,r)^\alpha}{d(s,r')^\alpha} + \frac{d(s,r)^\alpha}{d(s',r)^\alpha} \leq \frac{1}{2 \cdot 3^\alpha \cdot (4\beta + 2)}$*  **then**
  - 4          $\perp$  add  $\ell'$  to  $\mathcal{L}'_j$ ;
  - 5 compute power assignment  $\sigma$  such that all SINR constraints are fulfilled;
  - 6 return  $(\mathcal{L}'_1, \dots, \mathcal{L}'_k)$ ;
-

the channel he obtains. Hence, we will focus on  $S$  to denote simply *the set of bidders allocated to at least one channel* by the algorithm. Similarly, we use  $S^*$  to denote the set of bidders receiving at least one channel in the optimum.

**Case 1:**  $b(\mathcal{A}^*) < 8 \cdot B^*$  To bound the social welfare in this case, we only consider the event that  $SEC-PRICE = 1$ . This happens with probability  $\varepsilon$ . We know that, provided that  $SEC-PRICE = 1$ , the social welfare of the computed solution is precisely  $B^*$ . That is, the expected social welfare is at least

$$\mathbf{E}[b(\mathcal{A})] \geq \varepsilon \cdot B^* \geq \frac{\varepsilon}{8} \cdot b(\mathcal{A}^*) .$$

This completes the proof for the first case.

**Case 2:**  $b(\mathcal{A}^*) \geq 8 \cdot B^*$  Let  $i^*$  be a bidder with the highest bid  $b_{i^*} = B^*$ . To bound the expected revenue, we only consider cases, in which  $SEC-PRICE = 0$  and  $i^* \in \text{STAT}$ .

Under the assumption that  $i^* \in \text{STAT}$ , we denote by  $S_j$  the set of bidders getting at least one channel when  $X = j$ . In particular,  $S_j$  is the set of bidders allocated by the packing algorithm when executed on the set  $M_j = \{i \in \text{FIXED} \mid b_i \geq 2^{-j} \cdot B^*\}$ . Using  $p = 2^{-j} \cdot B^*$ , the expected revenue of the mechanism can be bounded by

$$\frac{1}{\lceil \log n \rceil + 2} \cdot \mathbf{E} \left[ \sum_{j=0}^{\lceil \log n \rceil + 1} 2^{-j} \cdot B^* \cdot |S_j| \mid i^* \in \text{STAT} \right] .$$

By linearity of expectation, this is equal to

$$\frac{1}{\lceil \log n \rceil + 2} \sum_{j=0}^{\lceil \log n \rceil + 1} 2^{-j} \cdot B^* \cdot \mathbf{E}[|S_j| \mid i^* \in \text{STAT}] . \quad (3)$$

Furthermore, let us define  $S_j^* = \{i \in S^* \mid b_i \geq 2^{-j} \cdot B^*\}$  as the set of bidders receiving a channel in the optimum and having a bid at least  $2^{-j} \cdot B^*$ . Obviously,  $S_j^* \cap \text{FIXED}$  is a subset of  $M_j$ , and it can be turned into a feasible allocation because it allocates only a subset of bidders from the optimum. Hence, using our assumption on the quality of algorithm **UnweightedPacking**, we have  $|S_j| \geq \kappa \cdot |S_j^* \cap \text{FIXED}|$  for a suitable constant  $\kappa$ . Furthermore, we can bound the expected size of  $S_j^* \cap \text{FIXED}$  by

$$\begin{aligned} \mathbf{E}[|S_j^* \cap \text{FIXED}| \mid i^* \in \text{STAT}] &= \mathbf{E}[|S_j^* \cap \text{FIXED} \setminus \{i^*\}|] \\ &= (1 - \varepsilon)|S_j^* \setminus \{i^*\}| \geq (1 - \varepsilon)(|S_j^*| - 1) . \end{aligned}$$

The value of this is at least

$$\frac{1}{\lceil \log n \rceil + 2} \sum_{j=0}^{\lceil \log n \rceil + 1} 2^{-j} B^* \kappa \cdot (1 - \varepsilon)(|S_j^*| - 1) \geq \frac{(1 - \varepsilon)\kappa}{\lceil \log n \rceil + 2} \left( \left( \sum_{j=0}^{\lceil \log n \rceil + 1} 2^{-j} B^* \cdot |S_j^*| \right) - 2B^* \right) .$$

To derive the approximation factor now, we have to bound  $b(\mathcal{A}^*)$  in terms of the cardinalities of the  $S_j^*$  sets. As a first step, we consider links that are not contained in any of these sets. For this purpose, let us define  $Y = \mathcal{N} \setminus S_{\lceil \log n \rceil + 1}^*$ . The contribution of these links to the social welfare is at most

$$\sum_{i \in Y} b_i \leq n \cdot 2^{-\lceil \log n \rceil - 1} \cdot B^* \leq \frac{b(\mathcal{A}^*)}{2} ,$$



because  $b(\mathcal{A}^*) \geq B^*$ . This yields for the remaining links

$$\sum_{j=0}^{\lceil \log n \rceil + 1} 2^{-j} \cdot B^* \cdot |S_j| \geq \frac{b(\mathcal{A}^*)}{4} .$$

Combining this bound with the previously obtained bound on revenue, we get that the expected revenue is at least

$$\frac{(1-\varepsilon)\kappa}{\lceil \log n \rceil + 2} \left( \frac{b(\mathcal{A}^*)}{4} - 2B^* \right) \geq \frac{(1-\varepsilon)\kappa}{\lceil \log n \rceil + 2} \cdot \frac{b(\mathcal{A}^*)}{8} .$$

For small  $\varepsilon$ , this shows that the revenue is in  $\Omega\left(\frac{\varepsilon \cdot \kappa}{\log n}\right) \cdot b(\mathcal{A}^*)$ . As no bidder has negative utility, this shows the same bound for the social welfare. With  $\kappa$  being constant, we obtain the claim of the theorem in Case 2 with  $SEC-PRICE = 0$  and  $i^* \in \text{STAT}$ .

Finally,  $SEC-PRICE = 0$  and  $i^* \in \text{STAT}$  happens with probability  $(1-\varepsilon) \cdot \varepsilon$ . In all other cases, we underestimate the welfare of the solution computed by the mechanism by 0. For the expected social welfare obtained by our algorithm this yields

$$\mathbf{E}[b(\mathcal{A})] \geq \frac{(1-\varepsilon)^2 \cdot \varepsilon \cdot \kappa}{(\lceil \log n \rceil + 2) \cdot 8} \cdot b(\mathcal{A}^*) .$$

This proves the claim for Case 2 and yields the theorem.  $\square$

## 4 General Technique for Universally Truthful Network Auctions

In this section, we demonstrate that our approach from the previous section allows to build universally truthful mechanisms for virtually all secondary spectrum auctions. To illustrate our results, we will focus on secondary network auctions studied Zhu et al. [27]. In this model, *secondary networks* are selfish users participating in an auction. Each user has a network  $G_i$  with a dedicated source node  $s_i$  and destination node  $d_i$  that are publicly known. User  $i$  wants to connect  $s_i$  to  $d_i$  without interference, and any interference-free  $s_i$ - $d_i$ -path has value  $w_i$  for her. The valuation  $w_i$  is private information. The allocation problem of the mechanism is to select a subset  $S \subseteq \mathcal{N}$  of the users. For each selected user  $i \in S$ , it has to allocate a path of links connecting  $s_i$  to  $d_i$  in  $G_i$ . For each used link in any of the  $G_i$  graphs, it furthermore has to select a channel from  $\mathcal{C}$ . The selection of paths and channels has to be free of interference. The interference constraint is modeled using an additional conflict graph  $H$  as follows. The vertex set of  $H$  consists of all links of the  $G_i$ -graphs. The edges of  $H$  indicate which links in the  $G_i$ -graphs collide due to interference. An allocation of channels to links and paths to bidders is only feasible if none of the used links that are on the same channel are connected by an edge in  $H$ .

The overall approach of the mechanism is as before. The users here are again single-parameter agents – their private information is the single numerical value  $w_i$  representing the value for an interference-free  $s_i$ - $d_i$ -path in  $G_i$ . Zhu et al [27] present a deterministic heuristic for the allocation problem that satisfies Myerson’s monotonicity condition [21] and can therefore be turned into a truthful mechanism. Furthermore, they discuss a truthful-in-expectation mechanism based on randomized meta-rounding with a provable approximation guarantee. The guarantee depends on parameters like the maximum length of an allocated path and structural properties of the conflict graph  $H$ .

---

**Algorithm 3:** Universally-truthful framework

---

```
1 For each bidder  $i$ : If there is no allocation that yields  $i$  as winner, remove  $i$  from  $\mathcal{N}$ ;  
2 Set  $SEC-PRICE$  to 1 with probability  $\varepsilon$ , otherwise to 0;  
3 if  $SEC-PRICE = 1$  then  
4   Let  $i^*$  be a bidder such that  $b_{i^*} = \max_{i \in \mathcal{N}} b_i$ ;  
5   Compute the allocation in which only  $i^*$  is a winner;  
6   Charge  $p_{i^*}(b) := \max_{i \in \mathcal{N} \setminus \{i^*\}} b_i$ ;  
7   All other bidders  $i$  are charged  $p_i(b) := 0$ ;  
8 else  
9   forall the  $i \in \mathcal{N}$  do  
10     $\perp$  add  $i$  to STAT with probability  $\varepsilon$ , otherwise to FIXED  
11    Set  $B := \max_{i \in \text{STAT}} b_i$ ;  
12    Choose  $X$  uniformly at random from  $\{0, 1, \dots, \lceil \log n \rceil + 1\}$ ;  
13    Set  $p := 2^{-X} \cdot B$ ;  
14    Set  $M := \{i \in \text{FIXED} \mid b_i \geq p\}$ ;  
15    Run UnweightedPacking( $M$ ) which returns allocation  $\mathcal{A}$ ;  
16    If  $i$  contained in  $\mathcal{A}$ , charge  $p_i(b) := p$ , otherwise  $p_i(b) := 0$ ;
```

---

There are further fundamental similarities with our single-hop scheduling approach. This scenario also exhibits a *packing structure*. Each bidder  $i$  either obtains some interference-free path, for which her valuation is  $w_i$ , or she is not served at all. In this sense, an allocation determines a set  $W$  of winners that become successful in the solution, we will say they are assigned a “license” for successful transmission. The social welfare of the allocation is given by  $\sum_{i \in W} w_i$ . By packing structure we refer to the property that for any feasible allocation, we can remove any winner  $i' \in W$  without changing the feasibility of the allocation for any other bidder. In both, secondary network and single-hop scheduling scenarios, this is due to the fact that by removing bidders, interference for the other bidders can only decrease. In the secondary network setting, removing bidders and their paths will never introduce new conflicts. The same is true for request scheduling, because the SINR can only increase by removing bidders and SINR constraints cannot be violated this way.

To formalize this property, we introduce a family  $\mathcal{W} \subseteq 2^{\mathcal{N}}$  consisting of sets of bidders. Each  $W \in \mathcal{W}$  represents a subset of bidders that are assigned a license in a particular feasible solution. That is, if we consider a fixed allocation, each  $i \in W$  obtains a license, whereas no  $i \notin W$  gets a license. If a scenario has a packing structure, the family  $\mathcal{W}$  is *downward closed*: For any  $W \in \mathcal{W}$  all subsets are contained in  $\mathcal{W}$  as well, i.e., if  $W \in \mathcal{W}$  and  $W' \subseteq W$ , then also  $W' \in \mathcal{W}$ .

For the single-hop scenario considered in Section 3,  $\mathcal{W}$  contains all sets  $W \subseteq \mathcal{N}$  of links that can be allocated to the channels in  $\mathcal{C}$  such that there is a power assignment fulfilling all SINR constraints. In case of the secondary network setting, it contains all sets  $W \subseteq \mathcal{N}$  for which there is an allocation to channels guaranteeing an interference-free selection of paths.

Based on this property, we can generalize the framework introduced in the previous section as shown in Algorithm 3. We first adjust the set of bidders  $\mathcal{N}$  to remove all bidders that can never be part of a feasible allocation. Because  $\mathcal{W}$  is downward closed, it suffices to check for each single  $i \in \mathcal{N}$  if there is a feasible allocation with  $W = \{i\}$ . This step is unnecessary for single-hop scheduling with power control in the previous section, as sufficiently large powers guarantee that every single bidder could become successful if she is the only one assigned to transmit.

In the main routine, our mechanism relies on an algorithm `UnweightedPacking`( $M$ ). Given a set  $M$  of candidate bidders, it calculates a feasible allocation  $\mathcal{A}$ , in which only the bidders in  $W$  are served. Here algorithm `UnweightedPacking` again neglects the bids.

**Theorem 3.** *Algorithm 3 is universally truthful.*

Proving truthfulness can be done by literally the same arguments as used in the proof of Theorem 1. Also, we can again show an approximation guarantee based on the applied algorithm `UnweightedPacking`( $M$ ). In the following, we assume that this algorithm is a  $\psi$ -approximation, meaning that for the winning set  $W$ , we have  $|W| \geq \psi \cdot \max_{W' \in \mathcal{W} \cap 2^M} |W'|$ . Given this guarantee, we can show the following approximation factor for the complete mechanism.

**Theorem 4.** *Let  $\mathcal{A}$  be the allocation returned by the mechanism and  $\mathcal{A}^*$  be the allocation optimizing social welfare. Then we have*

$$\mathbf{E}[b(\mathcal{A})] = \Omega\left(\frac{\varepsilon\psi}{\log n}\right) \cdot b(\mathcal{A}^*) .$$

The proof of this theorem is almost identical to the proof of Theorem 2 (with winner sets  $W$  taking the role of sets  $S$  of scheduled users) and is presented for completeness in the appendix.

## 5 Applications and Results

In the following, we discuss some applications of our approach, for which we obtain universally truthful mechanisms with non-trivial worst-case guarantees using our general technique.

**Power Control** As seen above, we can design universally truthful mechanisms for single-hop request scheduling in the SINR model with power control that yield approximation factors of  $O(\log n)$ . If we cannot use power control and must resort to fixed power schemes, there might exist request that do not overcome noise with their fixed power assignment. Here we can use the general framework and implement the allocation algorithm `UnweightedPacking`( $M$ ) by, e.g., adapting the algorithms of Goussevskaia et al. [10], Halldórsson and Mitra [13], or Halldórsson et al. [12]. While these algorithms have been designed to maximize the number of feasible requests on a single channel, it is rather easy to observe that they maintain their approximation guarantees even if we can assign requests to more than one channel (for completeness, see Theorem 5 in the Appendix). This way, we obtain universally truthful mechanisms for single-hop scheduling with, e.g., approximation factors of  $O(\log n)$  for uniform, linear, and square-root power assignments.

**Interference Models** Instead of SINR interference, we can also use any graph-based interference models, such as the protocol model or the disk-graph model. For the two latter models, there exist constant-factor approximation algorithms to maximize the number of feasible requests on a single channel. For the disc-graph model a 5-approximation algorithm is folklore. For the protocol model, one can construct conflict graphs with constant inductive independence number [16, 24], using which a simple greedy algorithm constructs a constant-factor approximation for capacity maximization [1]. As we prove in Theorem 5 in the Appendix, we can generically extend these algorithms to constant-factor approximation algorithms for capacity maximization with multiple channels. Using these algorithms for `UnweightedPacking`( $M, \mathcal{C}$ ), we obtain universally truthful mechanisms for single-hop scheduling in disc graph or protocol models with approximation factors of  $O(\log n)$ .

**Secondary Networks** In the secondary network model above, we can use algorithms described by Zhu et al [27] as a subroutine for `UnweightedPacking`( $M$ ) in our general framework. In this way, we obtain a universally truthful mechanism with approximation factor  $O(g(L_{\max}, \Delta) \cdot \log n)$ . For a description of the factor  $g(L_{\max}, \Delta)$ , see [27].

**General Packing Problems** More generally, *every* allocation problem in (wireless) networks that has a packing structure and allows a non-trivial approximation algorithm can be turned into a universally truthful mechanism using our approach, while spending only an additional  $O(\log n)$  factor in the guarantee. This includes, e.g., a large variety of routing, scheduling, congestion, and assignment problems in networks.

## 6 Conclusion

In this paper, we have presented a technique to design universally truthful secondary spectrum auctions that extends the literature on existing spectrum auctions in several ways. Furthermore, it allows to turn arbitrary approximation algorithms into truthful mechanisms in a large variety of previously studied models. To apply our technique to further settings, one only needs to adapt the respective allocation algorithm `UnweightedPacking`( $M$ ), which does not have to fulfill any further requirements. In any case, we guarantee that bidders cannot benefit from misreporting their valuation, even if they are risk-averse or risk-seeking.

An interesting topic for future work is to generalize the approach to multi-parameter mechanism design. In terms of secondary spectrum auctions, this means, for example, that each user’s valuation depends on the amount of interference that it is exposed to, rather than only considering whether interference is too high or not.

## References

- [1] Karhan Akcoglu, James Aspnes, Bhaskar DasGupta, and Ming-Yang Kao. Opportunity cost algorithms for combinatorial auctions. In E. Kontogiorgos, B. Rustem, and S. Siokos, editors, *Applied Optimization 74: Computational Methods in Decision-Making, Economics and Finance*, pages 455–479. Kluwer, 2002.
- [2] Matthew Andrews and Michael Dinitz. Maximizing capacity in arbitrary wireless networks in the SINR model: Complexity and game theory. In *Proc. 28th IEEE Conf. Computer Communications (INFOCOM)*, pages 1332–1340, 2009.
- [3] Randall Berry, Michael Honig, and Rakesh Vohra. Spectrum markets: Motivation, challenges, and implications. *IEEE Communications Magazine*, 2010.
- [4] Robert Carr and Santosh Vempala. Randomized metarounding. *Random Struct. Algorithms*, 20(3):343–352, 2002.
- [5] Edward Clarke. Multipart pricing of public goods. *Public Choice*, 8:17–33, 1971.
- [6] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Truthful randomized mechanisms for combinatorial auctions. *J. Comput. Syst. Sci.*, 78(1):15–25, 2012.

- [7] Thomas Erlebach, Klaus Jansen, and Eike Seidel. Polynomial-time approximation schemes for geometric graphs. *SIAM J. Comput.*, 34(6):1302–1323, 2005.
- [8] Ajay Gopinathan and Zongpeng Li. A prior-free revenue maximizing auction for secondary spectrum access. In *Proc. 30th IEEE Conf. Computer Communications (INFOCOM)*, pages 86–90, 2011.
- [9] Ajay Gopinathan, Zongpeng Li, and Chuan Wu. Strategyproof auctions for balancing social welfare and fairness in secondary spectrum markets. In *Proc. 30th IEEE Conf. Computer Communications (INFOCOM)*, pages 3020–3028, 2011.
- [10] Olga Goussevskaya, Magnús Halldórsson, Roger Wattenhofer, and Emo Welzl. Capacity of arbitrary wireless networks. In *Proc. 28th IEEE Conf. Computer Communications (INFOCOM)*, pages 1872–1880, 2009.
- [11] Theodore Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.
- [12] Magnús Halldórsson, Stephan Holzer, Pradipta Mitra, and Roger Wattenhofer. The power of non-uniform wireless power. In *Proc. 24th Symp. Discrete Algorithms (SODA)*, pages 1595–1606, 2013.
- [13] Magnús Halldórsson and Pradipta Mitra. Wireless capacity with oblivious power in general metrics. In *Proc. 22nd Symp. Discrete Algorithms (SODA)*, pages 1538–1548, 2011.
- [14] Magnús Halldórsson and Roger Wattenhofer. Wireless communication is in APX. In *Proc. 36th Intl. Coll. Automata, Languages and Programming (ICALP)*, volume 1, pages 525–536, 2009.
- [15] Martin Hoefer and Thomas Kesselheim. Secondary spectrum auctions for symmetric and submodular bidders. In *Proc. 13th Conf. Electronic Commerce (EC)*, pages 657–671, 2012.
- [16] Martin Hoefer, Thomas Kesselheim, and Berthold Vöcking. Approximation algorithms for secondary spectrum auctions. In *Proc. 23rd Symp. Parallelism in Algorithms and Architectures (SPAA)*, pages 177–186, 2011.
- [17] Thomas Kesselheim. A constant-factor approximation for wireless capacity maximization with power control in the SINR model. In *Proc. 22nd Symp. Discrete Algorithms (SODA)*, pages 1549–1559, 2011.
- [18] Thomas Kesselheim. Approximation algorithms for wireless link scheduling with flexible data rates. In *Proc. 20th European Symposium on Algorithms (ESA)*, pages 659–670, 2012.
- [19] Ron Lavi and Chaitanya Swamy. Truthful and near-optimal mechanism design via linear programming. *J. ACM*, 58(6):25, 2011.
- [20] Thomas Moscibroda and Roger Wattenhofer. The complexity of connectivity in wireless networks. In *Proc. 25th IEEE Conf. Computer Communications (INFOCOM)*, pages 1–13, 2006.
- [21] Roger Myerson. Optimal auction design. *Math. Oper. Res.*, 6:58–73, 1981.
- [22] Tim Nieberg, Johann Hurink, and Walter Kern. Approximation schemes for wireless networks. *ACM Trans. Algorithms*, 4(4), 2008.

- [23] William Vickrey. Counterspeculation, auctions and competitive sealed tenders. *J. Finance*, 16:8–37, 1961.
- [24] Peng-Jun Wan. Multiflows in multihop wireless networks. In *Proc. 10th Symp. Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 85–94, 2009.
- [25] Xia Zhou, Sorabh Gandhi, Subhash Suri, and Haitao Zheng. eBay in the Sky: Strategy-proof wireless spectrum auctions. In *Proc. 14th Intl. Conf. Mobile Computing and Networking (MOBICOM)*, pages 2–13, 2008.
- [26] Xia Zhou and Haitao Zheng. TRUST: A general framework for truthful double spectrum auctions. In *Proc. 28th IEEE Conf. Computer Communications (INFOCOM)*, pages 999–1007, 2009.
- [27] Yuefei Zhu, Baochun Li, and Zongpeng Li. Truthful spectrum auction design for secondary networks. In *Proc. 31st IEEE Conf. Computer Communications (INFOCOM)*, pages 873–881, 2012.

## A Appendix

### A.1 Missing Proofs

*Proof.* (of Theorem 4) We denote the highest bid by  $B^* = \max_{i \in \mathcal{N}} b_i$ . Let  $W$  denote the set of winning bidders in the computed allocation. Similarly, let  $W^*$  to denote the set of winning bidders in allocation that maximizes social welfare.

**Case 1:**  $b(\mathcal{A}^*) < 8 \cdot B^*$  To bound the social welfare in this case, we only consider the event that  $SEC-PRICE = 1$ . This happens with probability  $\varepsilon$ . We know that, provided that  $SEC-PRICE = 1$ , the social welfare of the computed solution is precisely  $B^*$ . That is, the expected social welfare is at least

$$\mathbf{E}[b(\mathcal{A})] \geq \varepsilon \cdot B^* \geq \frac{\varepsilon}{8} \cdot b(\mathcal{A}^*) .$$

This completes the proof for the first case.

**Case 2:**  $b(\mathcal{A}^*) \geq 8 \cdot B^*$  Let  $i^*$  be a bidder with the highest bid  $b_{i^*} = B^*$ . To bound the expected revenue, we only consider cases, in which  $SEC-PRICE = 0$  and  $i^* \in \text{STAT}$ .

Under the assumption that  $i^* \in \text{STAT}$ , we denote by  $W_j$  the set of winning bidders when  $X = j$ . That is,  $W_j$  is the set of bidders allocated by the packing algorithm when executed on the set  $M_j = \{i \in \text{FIXED} \mid b_i \geq 2^{-j} \cdot B^*\}$ . Using  $p = 2^{-j} \cdot B^*$ , the expected revenue of the mechanism can be bounded by

$$\frac{1}{\lceil \log n \rceil + 2} \cdot \mathbf{E} \left[ \sum_{j=0}^{\lceil \log n \rceil + 1} 2^{-j} \cdot B^* \cdot |W_j| \mid i^* \in \text{STAT} \right] .$$

By linearity of expectation, this is equal to

$$\frac{1}{\lceil \log n \rceil + 2} \sum_{j=0}^{\lceil \log n \rceil + 1} 2^{-j} \cdot B^* \cdot \mathbf{E}[|W_j| \mid i^* \in \text{STAT}] . \quad (4)$$

Furthermore, let us define  $W_j^* = \{i \in W^* \mid b_i \geq 2^{-j} \cdot B^*\}$  as the set of winning bidders in the optimum and having a bid at least  $2^{-j} \cdot B^*$ . Obviously,  $W_j^* \cap \text{FIXED}$  is a subset of  $M_j$ . As  $\mathcal{W}$  is downward closed, it is also contained in  $\mathcal{W}$ . Hence, using our assumption on the quality of algorithm **UnweightedPacking**, we have  $|W_j| \geq \psi \cdot |W_j^* \cap \text{FIXED}|$ . Furthermore, we can bound the expected size of  $W_j^* \cap \text{FIXED}$  by

$$\begin{aligned} & \mathbf{E}[|W_j^* \cap \text{FIXED}| \mid i^* \in \text{STAT}] \\ &= \mathbf{E}[|W_j^* \cap \text{FIXED} \setminus \{i^*\}|] \\ &= (1 - \varepsilon) |W_j^* \setminus \{i^*\}| \\ &\geq (1 - \varepsilon) (|W_j^*| - 1) . \end{aligned}$$

The value of this is at least

$$\begin{aligned} & \frac{1}{\lceil \log n \rceil + 2} \sum_{j=0}^{\lceil \log n \rceil + 1} 2^{-j} \cdot B^* \cdot \psi \cdot (1 - \varepsilon) (|W_j^*| - 1) \\ &\geq \frac{(1 - \varepsilon)\psi}{\lceil \log n \rceil + 2} \left( \left( \sum_{j=0}^{\lceil \log n \rceil + 1} 2^{-j} \cdot B^* \cdot |W_j^*| \right) - 2B^* \right) . \end{aligned}$$

To derive the approximation factor now, we have to bound  $b(\mathcal{A}^*)$  in terms of the cardinalities of the  $W_j^*$  sets. As a first step, we consider bidders that are not contained in any of these sets. For this purpose, let us define  $Y = \mathcal{N} \setminus S_{\lceil \log n \rceil + 1}$ . The contribution of these bidders to the social welfare is at most

$$\sum_{i \in Y} b_i \leq n \cdot 2^{-\lceil \log n \rceil - 1} \cdot B^* \leq \frac{b(\mathcal{A}^*)}{2} ,$$

because  $b(\mathcal{A}^*) \geq B^*$ . This yields for the remaining links

$$\sum_{j=0}^{\lceil \log n \rceil + 1} 2^{-j} \cdot B^* \cdot |W_j| \geq \frac{b(\mathcal{A}^*)}{4} .$$

Combining this bound with the previously obtained bound on revenue, we get that the expected revenue is at least

$$\frac{(1-\varepsilon)\psi}{\lceil \log n \rceil + 2} \left( \frac{b(\mathcal{A}^*)}{4} - 2B^* \right) \geq \frac{(1-\varepsilon)\psi}{\lceil \log n \rceil + 2} \cdot \frac{b(\mathcal{A}^*)}{8} .$$

For small  $\varepsilon$ , this shows that the revenue is in  $\Omega\left(\frac{\varepsilon \cdot \psi}{\log n}\right) \cdot b(\mathcal{A}^*)$ . As no bidder has negative utility, this shows the same bound for the social welfare.

Finally,  $SEC-PRICE = 0$  and  $i^* \in \text{STAT}$  happens with probability  $(1-\varepsilon) \cdot \varepsilon$ . In all other cases, we underestimate the welfare of the solution computed by the mechanism by 0. For the expected social welfare obtained by our algorithm this yields

$$\mathbf{E}[b(\mathcal{A})] \geq \frac{(1-\varepsilon)^2 \cdot \varepsilon \cdot \psi}{(\lceil \log n \rceil + 2) \cdot 8} \cdot b(\mathcal{A}^*) .$$

This proves the claim for case 2 and yields the theorem.  $\square$

## A.2 Extending Capacity Maximization Algorithms to Multiple Channels

In order to apply algorithms that are meant to solve the single-channel allocation problem in the context of multiple channels, we can use the following reduction: We iteratively fill each of the sets  $\mathcal{L}_1, \dots, \mathcal{L}_k$  by applying the single-channel algorithm on all remaining users. We can show that we only lose a constant-factor in terms of the approximation guarantee.

**Theorem 5.** *Given an algorithm with approximation factor of  $\psi$  for the unweighted single-channel allocation problem, we obtain a  $((1 - 1/e)\psi)$ -approximation algorithm for the unweighted  $k$ -channel problem.*

*Proof.* Let  $\mathcal{L}_1, \dots, \mathcal{L}_k$  be the set of users selected for the respective channels by the algorithm. Let  $\text{OPT}$  be the total number of selected users in the optimal solution.

Let us consider the  $j$ th run of the algorithm, which selects the users for set  $\mathcal{L}_j$ . We know that at this stage at most  $\sum_{i=1}^{j-1} |\mathcal{L}_i|$  have been selected in previous rounds. This means that at least  $\text{OPT} - \sum_{i=1}^{j-1} |\mathcal{L}_i|$  users that are selected in the optimal solution are still unselected. In the optimal allocation these are allocated to at most  $k$  different channels, meaning that at least  $\frac{\text{OPT} - \sum_{i=1}^{j-1} |\mathcal{L}_i|}{k}$  can share a channel. By using the approximation guarantee of the single-channel algorithm, we get

$$|\mathcal{L}_j| \geq \psi \frac{\text{OPT} - \sum_{i=1}^{j-1} |\mathcal{L}_i|}{k} .$$



Adding  $\sum_{i=1}^{j-1} |\mathcal{L}_i|$  to both sides, we get

$$\begin{aligned} \sum_{i=1}^j |\mathcal{L}_i| &\geq \sum_{i=1}^{j-1} |\mathcal{L}_i| + \psi \frac{\text{OPT} - \sum_{i=1}^{j-1} |\mathcal{L}_i|}{k} \\ &= \left(1 - \frac{\psi}{k}\right) \sum_{i=1}^{j-1} |\mathcal{L}_i| + \frac{\psi}{k} \text{OPT} . \end{aligned}$$

Solving this recursion, this yields

$$\begin{aligned} \sum_{i=1}^k |\mathcal{L}_i| &\geq \sum_{i=0}^{k-1} \left(1 - \frac{\psi}{k}\right)^i \cdot \frac{\psi}{k} \text{OPT} \\ &= \frac{\left(1 - \frac{\psi}{k}\right)^k}{1 - \left(1 - \frac{\psi}{k}\right)} \cdot \frac{\psi}{k} \text{OPT} \\ &= \left(1 - \frac{\psi}{k}\right)^k \text{OPT} \\ &\geq \left(1 - \frac{1}{e}\right) \psi \text{OPT} . \end{aligned}$$

This shows the claim. □