EXPERIMENTAL ANALYSIS OF THROUGHPUT MAXIMIZATION FOR COMBINATORIAL SPECTRUM AUCTIONS.

Hörður Ingi Björnsson
Master of Science
Computer Science
June 2013
School of Computer Science
Reykjavík University

M.Sc. PROJECT REPORT

ISSN 1670-8539

Experimental analysis of throughput maximization for combinatorial spectrum auctions.

by

Hörður Ingi Björnsson

Project report submitted to the School of Computer Science
at Reykjavík University in partial fulfillment of
the requirements for the degree of
Master of Science in **Computer Science**

June 2013

Project Report Committee:

Magnús Már Halldórsson, Supervisor
Professor, Computer Science

Eyjólfur Ingi Ásgeirsson
Assistant Professor, Science and Engineering

Ýmir Vigfússon
Assistant Professor, Computer Science

Experimental analysis of throughput maximization for combinatorial spectrum auctions.

Hörður Ingi Björnsson

June 2013

Abstract

A somewhat artificial scarcity of spectrum is due to static allocation to so called Primary Users (PUs) who use it only intermittently. A promising solution uses market approaches to redistribute limited time access rights to so called Secondary Users (SUs) (Hoefer, Kesselheim, & Vöcking, 2011). A reasonable approach, concisely termed "eBay in the Sky"(Zhou, Gandhi, Suri, & Zheng, 2008), auctions licenses for SUs regularly.

We employ the *physical* or *signal-to-interference-plus-noise-ratio* (SINR) model to model wireless communication in the plane. We auction m channels to n links randomly scattered in the plane, with probability p of attaching to a previously placed link, known as preferential attachment (Ásgeirsson et al., 2012). Links have symmetric and downward sloping valuations. Since finding an optimal allocation is known to be NP-hard (Goussevskaia, Oswald, & Wattenhofer, 2007) we compare several approximation algorithms in simulations, namely GREEDYLENGTH, GREEDYWEIGHT, BUCKETLENGTH, BUCKETWEIGHT, LOCALRATIO, GREEDYINTERFERENCE, and a Linear Programming Algorithm (LP) (Hoefer & Kesselheim, 2012). The performance of all algorithms is compared varying three dimensions, n, m and p. Algorithms GREEDYWEIGHT, LOCALRATIO and LP quite consistently rank first, second and third, respectively. In further simulations, varying valuations, LOCALRATIO outperformed GREEDYWEIGHT.

An attempt was made to turn GREEDYWEIGHT and LOCALRATIO into truthful mechanisms using VCG-like payments. Neither payment scheme turned out to be truthful. A known method turns LP solutions into truthful-in-expectation mechanisms (Lavi & Swamy, 2005). While there is no known truthful mechanism for GREEDYWEIGHT and LOCALRATIO, LP algorithm may be used at the price of lower social welfare. In practice GREEDYWEIGHT and LOCALRATIO may be used to allocate channels in a first-price-auction.

Greining með hermunum á hámörkun afkasta í fjölrása tíðniuppboðum.

Hörður Ingi Björnsson

Júní 2013

Útdráttur

Skortur á tíðni fyrir þráðlaus samskipti er að hluta tilkominn vegna þess að forgangsnotendur fá úthlutað tíðni til langs tíma en nota hana ekki stöðugt. Markaðsaðferðir gætu komið sér vel til að endurúthluta tíðni til skamms tíma til annarra notenda (Hoefer et al., 2011). Ein slík aðferð, nefnd "‘eBay á himni"’ (Zhou et al., 2008), býður reglulega upp skammtímanotkunarleyfi á tíðni til annarra notenda.

Við notumst við efnislega (SINR) líkanið fyrir þráðlaus samskipti á sléttu. Boðnar eru upp m rásir sem n tenglar geta boðið í. Með líkum p er tengill staðsettur nálægt einhverjum tengli sem er þegar á sléttunni, annars er staðsetning tengils valin af handahófi (Ásgeirsson et al., 2012). Tenglar hafa áhuga á því hversu margar rásir þeir fá, en bjóða minna fyrir hverja viðbótarrás. Samanlagt verðmat tengla fyrir rásir sem þeir fá úthlutað er skilgreint sem velferð. Sýnt hefur verið að ekki er hægt að finna kjörlausn á margliðutíma (Goussevskaia et al., 2007), því berum við saman velferð nokkurra nálgunarreiknita með hermunum, það er GREEDYLENGTH, GREEDYWEIGHT, BUCKETLENGTH, BUCKETWEIGHT, LOCALRATIO, GREEDYINTERFERENCE og reiknirit sem byggir á línulegri bestun (LP) (Hoefer & Kesselheim, 2012). Velferð er borin saman á þremur víddum, það er n, m og p. Reikniritin GREEDYWEIGHT, LOCALRATIO og LP, finna lausnir með mestri velferð. Verðmati tengla var breytt í frekari hermunum, þar var LOCALRATIO með meiri velferð en GREEDYWEIGHT.

Uppboð eru skilgreind sannorð ef bjóðendur hagnast á því að bjóða sitt rétta verðmat. Reynt var að búa til sannorð uppboð fyrir GREEDYWEIGHT og LOCALRATIO, með greiðslum sem svipa til VCG, en tókst í hvorugu tilviki. Vitað er um aðferð sem breytir lausnum úr línulegum bestunarlíkönum í væntanlega sannorð uppboð (Lavi & Swamy, 2005). Meðan ekki er vitað um sannorð uppboð fyrir GREEDYWEIGHT and LOCALRATIO, er hægt að notast við LP reiknitið með aðeins minni velferð. Notast má við GREEDYWEIGHT og LOCALRATIO, ef tenglar eru látnir borga sama verð og þeir bjóða.

Dedicated to my beloved wife Þórunn and son Benedikt Örn.

Acknowledgements

This research is supported by Rannsóknarsjóður research grant number 90032021.

First of all, I would like to thank my supervisor Dr. Magnús Már Halldórsson. He has given me guidance through my studies, a direction for my research and enough freedom to explore my own ideas. This work could not have been done without him.

I would like to thank Dr. Eyjólfur Ingi Ásgeirsson for numerous discussions and suggestions, which have enriched and improved this work immensely.

Special thanks go to Dr. Pradipta Mitra, who was always willing to take time off explaining to me the fundamentals of wireless communication.

I am also very grateful to other members of our discussion group, Dr. Ýmir Vigfússon and Dr. Henning Úlfarsson. Our meetings have broadened my horizon and if I have learned something about carrying out rigorous research, it is especially thanks to our meetings.

Last but not least I would like to thank my wife Þórunn Guðmundsdóttir for her support and encouragement and my new born son Benedikt Örn Harðarson whose smile brightens every day.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The first public demonstration of a wireless packet data network became operational in June 1971. It was developed at the University of Hawaii and got the name ALOHAnet, or simply ALOHA (Schwartz & Abramson, 2009). In its simplest form, Pure ALOHA, a link that has data to send, just sends it. If there is a collision, the link waits random time and tries sending again. The maximum throughput in Pure ALOHA is $\frac{1}{2e} \approx 18.4\%$(Abramson, 1970).

In slotted ALOHA each link l transmits with probability $\frac{1}{n}$ in any given time slot, where n is the number of links wanting to transmit. Given that the links know the number n and that transmission is successful if exactly one link transmits, the probability of successful transmission in each slot is $\frac{1}{e} \approx 36.8\%$ (Roberts, 1975). The introduction of discrete time slots doubled the throughput achieved with the ALOHA protocol.

An obvious strength of ALOHA is that it is fully distributed. If there was a central authority that could schedule the wireless communication of wireless links 100% throughput could easily be achieved by scheduling one link in each time slot. Recent development in wireless communication goes even further. If links are far apart they can transmit simultaneously without interrupting each others communication.

In a seminal paper, capacity of wireless networks is studied under two models of interference: the *protocol* model and the *physical* or *signal-to-interference-plus-noise* (SINR) model (Gupta & Kumar, 2000). Let us look at each separately.

A link l_v is defined as a sender-receiver pair (s_v, r_v), where s_v and r_v are points in the plane. Suppose link l_v transmits over a channel. In the protocol model the transmission of l_v is successful if

$$|s_u - r_v| \geq (1 + \Delta)|s_v - r_v|, \tag{1.1}$$

for every other sender s_u simultaneously transmitting over the same channel, where $(1 + \Delta)$ is the ratio between transmission range and interference range. If the quantity $\Delta > 0$ a sender s_u can interfere with the communication of l_v even though it is further away from r_v than s_v.

In the SINR model the transmission of l_v is successful if

$$\frac{\frac{P_v}{|s_v - r_v|^\alpha}}{\sum_{l_w \in S \setminus \{l_v\}} \frac{P_w}{|s_w - r_v|^\alpha} + N} \geq \beta, \tag{1.2}$$

where P_v is the transmitting power of l_v, N is the ambient noise in the system, $\alpha > 2$ is the path-loss-coefficient and β is the minimum SINR required for successful communication. If $\beta > 1$ the sender s_v needs to be closer to r_v than any other sender s_w transmitting with the same power.

A feasible set S is a set of links which can transmit simultaneously such that the communication of all links $l \in S$ is successful. The capacity problem can then be defined as finding a set S of maximum size. Finding a schedule $S_1, \ldots, S_m$ of minimum number of slots, where all $l \in S$ are scheduled, is known as the scheduling problem. Both problems have been studied under both models of interference (Gupta & Kumar, 2000).

An advantage of the protocol model is that it is easy to analyze and many known algorithms can be readily applied. The capacity problem in the protocol model is the same as finding a maximum independent set. As an example, a simple greedy algorithm that forms a maximal independent set by choosing the node with minimum degree and removing its neighbours, achieves an approximation ratio of $(\Delta + 2)/3$ on graphs with maximum degree Δ (Halldórsson & Radhakrishnan, 1997). The algorithm could then be applied on the remaining links until there are none left to find a $\mathcal{O}(\log(n))$-approximation to the scheduling problem.

The disadvantage of the protocol model is that it fails to take into account the additive effect of wireless signals. Although a link l_w far away from link l_v would not interfere with its transmission, the added interference of multiple far away links might. The SINR model on the other hand takes that into account and is therefore considered more realistic. We will only concern ourselves with the SINR model from here on. The disadvantage of the SINR model is that it is harder to analyze and interference among links cannot be represented by an unweighted conflict graph (Iyer, Rosenberg, & Karnik, 2006).

Plenty of algorithmic results exist for wireless communication in the SINR model, i.e. a simple greedy algorithm which finds a $\mathcal{O}(1)$-approximation to the capacity problem. The algorithm orders links in increasing order of length and greedily adds them to the set S as long as the affectance from the links in S does not exceed a constant c. Then outputs the feasible set S. Repeated application of this algorithm can then be used to find a $\mathcal{O}(\log(n))$-approximation to the scheduling problem (Goussievskaia, Halldorsson, Wattenhofer, & Welzl, 2009).

A natural extension to the capacity problem is giving links weights and maximize the sum of weights possible to schedule under SINR constraints, known as the weighted capacity problem. A simple bucket algorithm approximates the weighted capacity problem. Divide links into buckets, such that a link l_v has at most twice the weight of any other link l_u in the bucket. Then we apply the simple greedy algorithm on each bucket and return the allocation with the highest sum of weights. This achieves a $\mathcal{O}(\log(n))$-approximation to the weighted capacity problem. We will later see that spectrum auctions are an extension of the weighted capacity problem.

Spectrum management is a major challenge of today's wireless networks. We are observing a transition from voice-only communication to multimedia type applications. As a result there is a growing need for higher data rates. Wireless communication takes place in the natural frequency spectrum, which although in principle is infinite, only a finite part of it is usable for wireless communication. The current static allocation of frequency cannot support the increasing requirements (Yücek & Arslan, First Quarter 2009). One way of increasing throughput in wireless communication is allocating frequency dynamically to multiple non-interfering links.

Spectrum is becoming a scarce resource but this scarcity is somewhat artificial. Demands for services vary at different times and in different areas. Causing frequency bands licensed to one application to become overloaded while other bands are idle at the same time. A promising solution to this artificial scarcity is to use market approaches to redistribute access rights for a limited time (Hoefer et al., 2011).

In the cognitive radio literature, the owner of the rights to use a specific frequency band is known as *primary user*. Parts of the spectrum currently unused by primary users can be offered to so-called *secondary users* for use in a local area. A reasonable approach, concisely termed "eBay in the Sky"(Zhou et al., 2008), is to auction licenses for secondary users regularly. Primary users will get a share in the price paid for temporary utilization of their frequency band and will therefore have an incentive to report when and where it will not be in use.

This leads us to our setting, where we have n secondary users and m channels set up for auction. We assume bidders have submodular, symmetric valuations, and will design algorithms able to handle this kind of valuations specifically. The channel allocation problem is about finding an allocation of channels that maximizes the sum of valuations.

The valuation of each bidder will represent its weight, and the channel allocation problem becomes an extension to the weighted capacity problem. We now have m channels and want to maximize the sum of weights over all channels. Furthermore we have to deal with strategic behaviour of links through mechanism design.

In Chapter 2 we will explain the SINR model of interference further and introduce important concepts in the auction literature. We will generalize combinatorial auctions to auctions with weighted conflict graphs, such as SINR constraints and explain inductive independence number ρ of graphs, a non-standard graph parameter used in two of the algorithms.

In Chapter 3 we will go through several algorithms for the channel allocation problem. Some of them have no proven approximation guarantees, and some will even be shown to perform arbitrarily poorly on special cases. This does not mean that they will necessarily perform poorly in real situations or randomly generated instances. An attempt was made to design incentive compatible mechanisms for two of the algorithms. VCG-like mechanisms using the allocations of the algorithms instead of the optimal solution are proposed (see Section 2.2.2 for more on VCG mechanism).

In Chapter 4 we will present simulation results for the algorithms presented. We will test different values of ρ for our LP algorithm and compare the fractional solution to the rounded solution and an optimal solution. We find that the optimal solution to the channel allocation problem can be found in reasonable time for $n \leq 40$ and $m \leq 4$. We will then compare the performance of the algorithms presented, except optimal, on three dimensions, the number of links n, number of channels m and the density of links, modeled as probability p of links attaching to previously generated links. We will compare the performance of two algorithms with payment computation. At the end of the chapter we will further compare greedyWeight and localRatio with differently generated valuations.

In Chapter 5 we will evaluate our simulation results and what they mean to an auctioneer. We provide some ideas for future work and discuss limitations of our simulations. Finally we give a conclusion and recommendation for an auctioneer conducting a spectrum auction based on our findings.

1.1 Related Work

The throughput maximization, or the capacity problem, has been shown to be NP-hard in graph based models (Jain, Padhye, Padmanabhan, & Qiu, 2003). The capacity problem in the SINR model has to take into account additive effects of wireless signals from all other simultaneously transmitting links. The capacity problem and the scheduling problem in the SINR model have also been shown to be NP-hard (Goussevskaia et al., 2007). For sake of efficiency we must settle for approximating the optimal solution. Several algorithms are known that approximate the capacity problem and the scheduling problem. We will go through a few of them since the capacity problem relates to spectrum auctions.

Algorithm GreedyPhysical orders links based on their interference number, defined in the following way. Interference number of link $l_v \in L$ is the number of links $l_u \in L$ prevented from communicating at the same time as l_v, even though no other links are transmitting. Links are then ordered in decreasing order of interference number and greedily added to the transmitting set S, if SINR feasibility conditions hold. The algorithm is once to approximate the capacity problem, but repeatedly until all links have been scheduled to approximate the scheduling problem (Brar, Blough, & Santi, 2006).

Algorithm ApproxDiversity divides the plane into sufficiently large squares and 4-colors the squares. To approximate the capacity problem, for each color i pick one link from each square of that color and put into the set C_i. Let C_i of maximum size be the set S of transmitting links. To approximate the weighted capacity ApproxDiversity, each color i picks the link with the highest weight in each square and puts into the set C_i. Then lets C_i of maximum weight to be the set S of transmitting links. ApproxDiversity approximates the scheduling problem by applying the algorithm for the unweighted capacity problem repeatedly until there are no links left to schedule (Goussevskaia et al., 2007).

The Algorithm ApproxLogN produces a $\mathcal{O}(1)$-approximation to the capacity problem and a $\mathcal{O}(\log(n))$-approximation to the scheduling problem. To approximate the capacity problem, the algorithm simply goes through the links in increasing order of length and greedily schedules links if the SINR feasibility conditions hold. To approximate the scheduling problem, the algorithm for the capacity problem is applied repeatedly until there are no links left to schedule (Goussievskaia et al., 2009).

In a paper by Goussievskaia et al. (2009), the performance of GreedyPhysical, Approx-Diversity and ApproxLogN was compared in simulations, both for random and clustered topology. The clustered topology aims to simulate heterogeneous density distribution, which is more typical in practice than uniform random distribution. For example, in a sensor network, some areas of interest have higher density of sensors whereas in other locations there are only a minimum amount of nodes to maintain connectivity. In the random topology GreedyPhysical presents slightly better performance than ApproxLogN for low density (less than 1600 nodes) but for high density, ApproxLogN produces on average 50% shorter schedules. ApproxLogN produces 2.5 times shorter schedules than ApproxDiversity in high density situations (25600 nodes). GreedyPhysical is not able to deal with clustered topology and computes 3 times longer schedule than ApproxLogN for low density scenario (100 nodes) and 60 times longer schedule for high density scenario (25600 nodes). The effects of varying the path loss exponent α was also analyzed. GreedyPhysical is invariant to changes in α. For $\alpha < 3$, GreedyPhysical outperforms the other two algorithms in the random topology. Its performance in the clustered topology is very poor for low values of α and deteriorates ApproxDiversity and ApproxLogN with increasing α. ApproxLogN produces the shortest schedule $\alpha \geq 3$ in random topology and for all values of α in the clustered topology (Goussievskaia et al., 2009). Simulation results imply that the topology of the network and the value of the path loss coefficient α affects the relative performance of the algorithms.

In a spectrum auction there may be multiple winners of each channel, subject to SINR constraints. Algorithms for the capacity problem may give us ideas and valuable insight into spectrum auctions. Although widely employed in other situations, the VCG mechanism has serious drawbacks in multi-winner auctions. It is known to produce low revenue and is vulnerable to bidder collusion. An auction framework, for spectrum auction in the unit disk graph model, has been proposed to meet these drawbacks using Nash bargaining solution to improve revenue and prevent collusion. The proposed mechanisms increase revenue by 15% to 30% depending on the defined radius of interference (Wu, Wang, Liu, & Clancy, 2008). It is uncertain whether these mechanisms could be extended to the more realistic physical model.

Our goal will be to maximize the social welfare obtained, but in many situations revenue maximization is a more natural goal. A primary user is often a privately owned company which may have paid a considerable amount for spectrum rights, may be reluctant to share its underutilized spectrum if revenue is too low. A suboptimal allocation mechanism, which runs in polynomial time has been proposed and shown in simulations to produce stable expected revenue (Jia, Zhang, Zhang, & Liu, 2009). However it is based on a cellular topology and does not employ the SINR model, thus it fails to take into account the additive effects of interference.

A natural extension to our work, is to apply auction mechanisms to routing. An optimal algorithm exists for the unicast routing problem. Ad-hoc VCG is a truthful and cost-efficient routing protocol for SINR networks with selfish agents. Finding a shortest route, also known as the unicast routing problem, is equivalent to the shortest path problem and can easily be solved in polynomial time. The protocol first finds a shortest route, and then calculates the external cost of each node, which then becomes its payment. Overpayment is defined as $\frac{\text{VCG}}{|\text{SP}|}$, where $|\text{SP}|$ is the cost of the shortest path and $\text{VCG} = |\text{SP}| + \text{P}$, where P is the total payment. The overpayment produced by ad-hoc VCG has an upper bound of $\frac{\text{VCG}}{|\text{SP}|} \leq 2^{\alpha+1}\frac{c_{max}}{c_{min}}$, where α is the path loss coefficient and $\frac{c_{max}}{c_{min}}$ is the ratio between the highest valuation and the lowest valuation of any node. Simulation results imply that over payment ratio can be expected to be well within this upper bound (Anderegg & Eidenbenz, 2003).

The multicast routing problem is known to be NP-hard, by equivalence to the Steiner tree problem. A few sub-optimal truthful mechanisms have been proposed and compared in simulations. Least cost path tree (LCPT), finds a unicast route from source to each of the destinations and combines these routes into a tree. Then VCG-like payments are applied for the least cost path. Pruning minimum spanning tree (PMST), constructs a minimum spanning tree and then prunes off edges that do not lead to a receiver. VCG like payments are then applied for the minimum spanning tree. Link weighted Steiner tree (LST), constructs an approximate minimum cost Steiner tree, denoted as LST(d). VCG-like payments are then computed using LST(d) for each edge in the network. The payments used are shown to be truthful and the performance of the algorithms are compared in simulations. The simulations showed that LST produced considerably less costly solutions than LCPT and PMST, with far less payments as well. The overpayment ratio was the highest for PMST, but similar for LCPT and LST (Wang, Li, & Wang, 2004). Although routing is not directly related to our work, it is an interesting extension.

The weighted maximum independent set problem (WMIS), is about finding an independent set of nodes in a graph, with the maximum sum of weights. Spectrum auctions are an extension of WMIS with conflict between links represented by weighted edges and care must be taken to handle strategic behaviour of bidders. The maximum independent set problem may be considered as a special case of WMIS where all nodes have equal weights. A non-standard graph parameter, inductive independence number ρ (further explained in Section 2.3), can be used to achieve $\mathcal{O}(\rho)$-approximation to the weighted maximum independent set problem (Ye & Borodin, 2009).

More closely related to our work is a linear programming (LP) formulation combined with probabilistic rounding which finds a $\mathcal{O}(\sqrt{k})$-approximation to combinatorial spectrum auctions where k is the number of channels, and a $\mathcal{O}(\sqrt{k}\log^2(n))$-approximation for wireless links in the SINR model. The linear programming phase utilizes the inductive independence ordering and the notion of backward neighbourhood. It can handle arbitrary valuations by querying demand oracles (Hoefer et al., 2011). The algorithm was further developed for symmetric and submodular bidders to achieve a $\mathcal{O}(\rho \log(n))$-approximation to the optimal social welfare in the SINR model (Hoefer & Kesselheim, 2012). Their approach can be combined with a LP-based framework to turn it into a truthful-in-expectation mechanism (Lavi & Swamy, 2005).

Chapter 2

Preliminaries

Before we can begin the coverage of the algorithms, there are some important definitions that need to be explained. First we will cover the physical interference model of wireless communication, which captures some fundamental characteristics of wireless communication, such as the additivity of signal strength and robustness of short links. Next the simple auctions, where there is one item and multiple bidders, will be introduced and important concepts in the auction literature will be explained. As we get a feel for the terminology used we move onto combinatorial auctions, where there are multiple items and multiple bidders, and each bidder can present different preferences for each subset of item. We explain how interference among wireless links may be represented by weighted conflict graphs. A realistic valuation function for wireless channels is presented and after that we will introduce a non-standard graph parameter, inductive independence number, which will be used in one of the algorithms covered. At the end of the chapter the channel allocation problem is formally stated.

2.1 Physical Interference Model

To capture inherent characteristics in wireless networks, it is important to choose the communication model well. Much existing literature concerns graph-based models, such as the protocol model and unit disk graphs, where interference is modeled as binary constraints based on a local measure. Nodes represent communication requests and an edge is put between two nodes if they are too close to each other. Such models may serve as a useful abstraction but have some limitations. One limitation is that they ignore the additivity of interference. A single transmitter that is far away may cause little interference, but the accumulated interference of several such transmitters may be enough to corrupt a transmission (Goussievskaia et al., 2009). Scheduling in the graph-based model usually involves finding a maximum independent set, matching or coloring and has been widely studied (Joo, Lin, & Shroff, 2008), (Kumar, Marathe, Parthasarathy, & Srinivasan, 2005a), (Kumar, Marathe, Parthasarathy, & Srinivasan, 2005b).

The physical model of wireless communication is a more realistic model. A signal is received successfully if the signal to interference plus noise ratio (SINR) is below a certain threshold. The interference measured at each receiver is the sum of the interference caused by all transmitters concurrently transmitting plus the ambient noise. The SINR model accounts for interference generated by transmitters located far away (Goussievskaia et al., 2009). Determining the capacity of networks under SINR constraints theoretically began with the pioneering work of Gupta and Kumar (Gupta & Kumar, 2000) and has been of much interest since.

Inefficiency of graph-based scheduling protocols in the SINR model is well documented and has been shown both theoretically and experimentally (Gronkvist & Hansson, 2001), (Moscibroda, Wattenhofer, & Weber, 2006). We employ the SINR model as it is more realistic and can account for transmitters far away.

2.1.1 The Model

A link l_v consists of a sender-receiver pair (s_v, r_v). We denote the Euclidean distance between two points p and q as $d(p, q)$ and the length of a link l_v is denoted $d_{vv} = d(s_v, r_v)$. The distance from a sender s_v to a receiver r_w is denoted as $d_{vw} = d(s_v, r_w)$.

A link l_v transmits with power P_v. A power assignment P is non-decreasing if $P_v \geq P_w$, when $d_{vv} \geq d_{ww}$ and sub-linear if $\frac{P_v}{d_{vv}^\alpha} \leq \frac{P_w}{d_{ww}^\alpha}$, when $d_{vv} \geq d_{ww}$ and α is the path loss exponent. Examples include *uniform* power assignment, where all links transmit with the same power, *linear* power assignment where $P_v = d_{vv}^\alpha$, and *mean* power assignment where $P_v = d_{vv}^{\alpha/2}$.

We assume the path loss radio propagation model for the reception of signals, where the received signal from transmitter w at receiver v is $P_{wv} = P_w/d_{wv}^\alpha$ and $\alpha > 2$.

Definition 1. *The physical interference model states that a receiver r_v successfully receives a message from s_v if and only if*

$$\frac{P_{vv}}{\sum_{l_w \in S \setminus \{l_v\}} P_{wv} + N} \geq \beta, \tag{2.1}$$

where N is the ambient noise, β denotes the minimum SINR required for successful communication, and S is the set of concurrently scheduled links.

If inequality 2.1 holds for each link $l \in S$, we say that S is feasible. Now we can define two fundamental problems which are closely related to secondary spectrum auction. We are given a set L of n links. The former is the capacity problem, where we want to find a feasible set $S \in L$ of links of maximum size. The latter is the scheduling problem, where we want to schedule all the links in L in the fewest possible number of time slots. Both problems can be solved approximately by a simple greedy algorithm which achieves a $\mathcal{O}(1)$-approximation for the first problem, and a $\mathcal{O}(\log n)$-approximation for the second problem (Goussievskaia et al., 2009).

More relevant to our concern though is the weighted capacity problem, where each link has a weight and we want to find a feasible set S which maximizes the sum of the weights of the links in the set. When links use linear power there exists a $\mathcal{O}(1)$-algorithm for the weighted capacity problem (Halldorsson & Mitra, 2011).

2.2 Auctions

2.2.1 Single Item Auctions

Setting up an auction is an efficient way to find the right price for an item when the market price is unknown. For the allocation of a single item there are traditionally four types of auctions (Krishna, 2002):

- *Open ascending bid auction (English auction)* in which the price is steadily raised by the auctioneer and bidders drop out once the price becomes too high. When there is only one bidder left he wins and pays the current price. A more common variant of this type of auction is where bidders make open bids which must be higher than previous bids until no bidder is willing to make a higher bid.
- *Open descending bid auction (Dutch auction)* in which the price starts with a sufficiently high price that no bidder is willing to pay. The price is then progressively lowered until one bidder is willing to pay the current price. He wins the auction and pays the current price.
- *First price sealed bid auction* in which bidders submit bids in a sealed envelope to the auctioneer. The auctioneer opens the envelopes and the highest bidder wins and pays the amount that he bids.

- *Second price sealed bid auction (Vickrey auction)* in which bidders submit bids in a sealed envelope to the auctioneer. The auctioneer opens the envelopes and the highest bidder wins and pays the amount of the second highest bid.

The *utility* of a bidder i is defined $u_i = b_i - p$, where b_i is the valuation of the allocation and p is the price he pays. If bidder i loses the auction, $b_i = 0$ and $p = 0$, resulting in utility $u_i = 0$. If an auction mechanism never pays anything to bidders, i.e. p is nonnegative, the mechanism has no positive transfers. If bidding truthfully never results in a negative utility u the mechanism is individually rational.

When bidders bid their true valuation of an item, they are said to bid *truthfully*. The Vickrey auction is commonly favoured since bidding truthfully is a weakly dominating strategy (Vickrey, 1961), meaning that no other strategy gives higher utility.

Theorem 1. *In a Vickrey auction a bidder i maximizes his utility u_i by bidding his true valuation b_i (Vickrey, 1961; Clarke, 1971; Groves, 1973).*

Proof. Assume that bidder i bids b_i^* instead of his true valuation b_i. Assume b_i is a winning bid; then any bid b_i^* which is also a winning bid results in the same payment p and hence the same utility u_i. Now if b_i^* is a losing bid, bidding b_i^* will result in $u_i^* = 0$ making this deviation not profitable. If we assume b_i is a losing bid, then any bid b_i^* which is also a losing bid will result in the same utility of zero. Any winning bid b_i^* will result in a payment higher than i's true valuation, hence resulting in negative utility. Thus we have shown that i cannot benefit from making a bid b_i^* different from his true valuation. $\square$

It is easy to see that the same does not hold true for the first price sealed bid auction, since bidding truthfully will always result in $u_i = 0$ whereas making a bid $b_i^* < b_i$ might give i positive utility. A rational bidder will therefore never bid his true valuation in a first price sealed bid auction.

A common goal of the auctioneer is to allocate the auctioned item to the bidder that will make most use of it. We assume that the valuation of each bidder reveals how able they are to use the auctioned item. We define *social welfare* in a single item auction as the true valuation of the winner of the auction. A single item auction is *socially efficient* if the winner of the auction is the one that values the item the most. In many applications, revenue maximization is a more natural goal of the auctioneer, however less is known about how to achieve that goal.

2.2.2 Combinatorial Auctions

In many situations it may be beneficial to auction multiple items at the same time, allowing bidders to submit bids for different subsets of items. Formally, there is a set of m indivisible items which are concurrently auctioned among n bidders. Every bidder i has a valuation function b_i which describes his preference for different subsets of items.

Definition 2. *A valuation b is a real-valued function that for each subset S of items, $b(S)$ is the value that bidder i obtains if he receives this subset of items. Here, a valuation must have "free disposal", i.e., be monotone: for $S \subseteq T$ we have that $b(S) \leq b(T)$, and it should be "normalized": $b(\emptyset) = 0$.*

The purpose of auctioning the items simultaneously is that a bidders valuation for a subset of items need not be equal to the sum of the valuations of the items within this subset. Implicit in the definition are two assumptions about bidder preferences: first, if bidder i wins bundle S and pays a price p his utility is $u_i = b_i(S) - p$. Second, we assume that bidders do not care how items they do not receive are allocated.

Definition 3. *An allocation of the items among the bidders is $S_1, \ldots, S_n$ where $S_i \cap S_j = \emptyset$ for every $i \neq j$. The social welfare obtained by an allocation is $\sum_i b_i(S_i)$. A socially efficient allocation is an allocation with maximum social welfare among all allocations.*

Usually the valuation function b_i of bidder i is private information, only known to i. The goal is to design an auction mechanism that finds a socially efficient allocation. There are several challenges we must deal with (Nisan, Roughgarden, Tardos, & Vazirani, 2007):

- *Computational complexity:* Even for simple cases the allocation problem is NP-complete.
- *Representation and communication:* Since bidders can have preferences for all different subsets, the valuation function of each bidder is exponential in size. How can we represent them and transfer enough information to the auctioneer?
- *Strategic behaviour:* Can we design a mechanism that eliminates strategic behaviour of bidders?

As in the single item auction we want to find a socially efficient allocation. To be able to do so we must ensure that bidders bid their true valuations. The *Vickrey-Clarke-Groves (VCG)* mechanism is known to induce truthful bidding. For a set of auctioned items $M = \{t_1, \ldots, t_m\}$ and a set of bidders $N = \{v_1, \ldots, v_n\}$, let B_N^M be a socially efficient allocation for a given bid combination (assuming truthful bids). A bidder v_i who wins a set of items $S \subseteq M$, pays $p_i = B^M_{N\setminus\{v_i\}} - B^{M\setminus S}_{N\setminus\{v_i\}}$, which is exactly the amount v_i needed to bid to win the auction. Hence, bidders maximize their utility by bidding truthfully. VCG mechanism is named after authors of papers that successively generalized the idea (Vickrey, 1961; Clarke, 1971; Groves, 1973).

However, VCG auction relies on finding a socially efficient allocation, which is analogous to finding a maximum independent set in a graph. Finding an independent set is known to be NP-complete, making VCG impractical.

Since finding the optimal solution is computationally intractable, we must settle for an approximation. Care must be taken to make approximation mechanisms truthful. Computationally tractable truthful mechanisms are known for some special cases of combinatorial auction, e.g. single minded bidders (Lehmann, O'Callaghan, & Shoham, 2002). A mechanism is truthful in expectation if a bidder maximizes his expected utility by bidding his true valuation. This means that on average a bidder will be best off by bidding truthfully. A risk-neutral bidder will therefore bid his true valuation. A method is known that turns linear programming algorithms into mechanisms that are truthful-in-expectation. The main idea is to apply VCG payments for the fractional solution and since we are in fractional domain, we can always scale down both the optimal LP solution and the VCG prices by α, which does not affect truthfulness (Lavi & Swamy, 2005).

2.2.3 Combinatorial Auction with SINR Conflict Graph

Conflict between two bidders can be represented by a conflict graph $C = \{V, E\}$, where V is the set of vertices and E is the set of edges. Vertices represent bidders and edges represent conflict between bidders (Hoefer et al., 2011). In a combinatorial auction with indivisible items the conflict graph would be a complete graph with unweighted edges.

Interference between wireless links in the SINR model would typically be represented by a weighted bidirectional conflict graph. We denote the edge from link l_v to l_w as C_{vw}. The weight of the edge $C_{vw} = \frac{P_{wv}}{P_{vv}}$ represents interference caused by sender s_w on receiver r_v. A feasible set S is one in which $\sum_{l_w \in S \setminus l_v} C_{vw} + \frac{N}{P_{vv}} \leq \frac{1}{\beta}$ for all $l_v \in S$, which is equivalent to the condition in Equation 2.1.

In a combinatorial auction $m \geq 1$ items are set up for auction among n bidders, in our case wireless channels and links, respectively. In a regular combinatorial auction there is only one winner for each item auctioned. However in a combinatorial auction with SINR conflict graph the winners of each channel may be any subset of links S for which SINR feasibility conditions hold.

2.2.4 Realistic Valuation Functions

In a combinatorial auction with n bidders and m items, there are 2^m different subsets of items. If we allow arbitrary valuations the number of valuations grows exponentially with the number of items auctioned. All these valuations quickly become hard to represent and work with as m grows. To deal with the complexities of arbitrary valuations we focus on a special case of valuation functions that may be reasonable to expect.

Given that the channels set up for auction are of similar wavelengths, they will have similar characteristics, such as transfer speed, distance we are able to transmit, and tolerance to obstacles. Therefore it is reasonable to assume that bidders will only be interested in the number of channels they get and not care about which specific channels; such valuations are called symmetric.

It may also be reasonable to assume that bidders are willing to pay less for each additional channel they receive than the one before, since they can order the communication requests by priority; such valuations are called downward sloping.

To make our setting more manageable we will assume that valuations are symmetric and downward sloping.

2.3 Inductive Independence Number

To improve the performance of graph algorithms it may be useful to analyze the graphs and see if they have some distinguishable properties. One of the algorithms we are going to investigate has a non-standard graph parameter as input.

Definition 4 (Inductive independence number ρ)**.** *An independent set is a set of vertices in a graph, no two of which are adjacent. 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 $|M \cap \{u \in V | \{u, v\} \in E, \pi(u) < \pi(v)\}| \leq \rho$.*

This means that for every vertex $v \in V$, the size of an independent set in the backward neighbourhood of v, that is the set of neighbours u of v with $\pi(u) < \pi(v)$, is at most ρ. Conflict graphs derived from many simple binary models of wireless communication, such as the protocol model and disk graphs have $\rho = \mathcal{O}(1)$ (Wan, 2009). In the case of spectrum auctions in the SINR model we have weighted conflict graphs with weighted independent sets as defined in Section 2.2.3. The inductive independence number of conflict graphs in the SINR model was shown to be $\rho = \mathcal{O}(\log n)$ (Hoefer & Kesselheim, 2012), but has recently been improved to $\rho = \mathcal{O}(1)$, except for uniform power(Halldórsson, Holzer, Mitra, & Wattenhofer, 2012).

2.4 Problem Statement

We are given a set of m channels that are going up for auction among n wireless links. We know that multiple wireless links can operate on the same channel simultaneously given that SINR constraints of Equation 2.1 hold. Our goal is to allocate the channels such that social welfare is maximized. This means that the channels should be allocated to the subsets of links that value them the most.

Problem 1. *Given a set of m channels and n links, where each link has symmetric and downward sloping valuations b_i. Design an auction mechanism that allocates the channels to the links $S_1, \ldots, S_n$, that maximizes social welfare,* $\max \sum b_i(|S_i|)$.

One such mechanism is VCG. However it relies on finding the socially efficient allocation, which is known to be computationally intractable except for small cases. For larger cases we must settle for approximation mechanisms. It can be shown that the channel allocation problem is NP-hard by equivalence to the maximum independent set.

Chapter 3

Algorithms

We present several algorithms aimed at approximating a socially efficient allocation, each explained in a separate subchapter. The purpose of this paper is not to prove theoretical bounds on performance, but to find algorithms that perform well in practice. Some of the algorithms presented will perform arbitrarily badly on special cases that are unlikely to come up in practice, but may do well in more realistic situations.

3.1 Algorithm GREEDYLENGTH

The main strength of GREEDYLENGTH is its simplicity. It is inspired by a simple greedy algorithm which achieves a $\mathcal{O}(1)$-approximation for the unweighted capacity problem and a $\mathcal{O}(\log n)$-approximation for the unweighted scheduling problem (Goussievskaia et al., 2009). Its worst case performance can be arbitrarily bad in the weighted case.

Algorithm 1: Algorithm GREEDYLENGTH for downward sloping symmetric valuations.

Order links in increasing order of length, $d^*_{11} \leq d^*_{22} \leq \cdots \leq d^*_{nn}$
$S \leftarrow \emptyset$
for $u = 1$ *to* n **do**
    if $\frac{P_{vv}}{\sum_{w \in S \cup \{l^*_u\} \setminus \{l_v\}} P_{wv} + N} \geq \beta, \forall v \in S \cup \{l^*_u\}$ **then**
        $S \leftarrow S \cup \{l^*_u\}$
    end
end

Output:

- Allocate all m channels to the links in set S.

The algorithm GREEDYLENGTH goes as follows. Start with an empty set S. Consider each link l_v in increasing order of length. If feasibility conditions hold for all links in $S \cup l_v$, then add l_v to S. When all links have been considered allocate all channels to the links in S (see Algorithm 1).

In our setting each link wants as many channels as it can get. Since the ordering of the links in GREEDYLENGTH does not depend on the valuation for an additional channel it is sufficient to run the algorithm once.

Algorithm 1 can perform arbitrarily badly on special cases. Consider the case where we have two links l_1 and l_2, of which l_1 is shorter and both cannot be scheduled simultaneously. Their valuations are b_1 and b_2, arbitrarily small and arbitrarily large respectively. One channel is set up for auction. Algorithm 1 would allocate the channel to l_1 and be arbitrarily far off from the optimal solution of allocating the channel to l_2.

3.2 Algorithm GREEDYWEIGHT

Algorithm 2, called GREEDYWEIGHT, allocates channels greedily in decreasing order of valuation. We first initialize empty sets $S_1, \dots, S_m$, one for each channel. Then for each channel k order the links in decreasing order of valuation for an additional channel and greedily add links to S_k if feasibility conditions hold.

Algorithm 2: Algorithm GREEDYWEIGHT for downward sloping symmetric valuations.

$S_1, \dots, S_m \leftarrow \emptyset$

for $k \in \{1, \dots, m\}$ **do**

  Order links in decreasing order of valuations for an additional channel, $b_1^* \geq b_2^* \geq \cdots \geq b_n^*$

  for $u = 1$ *to* n **do**

    if $\frac{P_{vv}}{\sum_{w \in S_k \cup \{l_u^*\} \setminus \{l_v\}} P_{wv} + N} \geq \beta, \forall v \in S_k \cup \{l_u^*\}$ **then**

      $S_k \leftarrow S_k \cup \{l_u^*\}$

    end

  end

end

Output:

- Allocate each channel i to set S_i of links, $i = 1, \dots, m$.

Notice that in GREEDYWEIGHT we need to reorder the links and find an allocation for each channel, where as in GREEDYLENGTH it was enough to find one allocation. This is due to that the valuation of links for an additional channel may be different if they were allocated a channel previously, but the length of the links stays the same.

It seems intuitive and fair to give links precedence based on their valuations. However you might have the situation where a long link excludes many other links from consideration because of a slightly higher valuation. Algorithm 2 can perform arbitrarily badly in such cases. Consider a long link l_1 with valuation b_1, which excludes short links $l_2, \ldots, l_n$ which all form a feasible set and each has valuation $b_1 - \epsilon$, where ϵ is an arbitrarily small number. One channel is set up for auction and Algorithm 2 allocates the channel to the long link. The number of links n can be arbitrarily large and the allocation of Algorithm 2 arbitrarily far off from the optimal allocation.

3.2.1 Payment Computation for GREEDYWEIGHT

Our payment scheme is inspired by the payment scheme of the VCG mechanism. In the VCG mechanism we go through each bidder and calculate its opportunity cost and that becomes its payment. The VCG mechanism relies on finding the optimal allocation (Vickrey, 1961; Clarke, 1971; Groves, 1973), which is probably seldom the case in Algorithm 2.

The payment computation for GREEDYWEIGHT goes as follows. Let $S \in L$ be the allocation found by GREEDYWEIGHT for an auction of m channels for the set of links L. Let $P_{n\times 1} = 0$ be a vector used to store the payment of each link. For all $v \in S$ in any order, let T_v be the allocation found by GREEDYWEIGHT for links $L \setminus \{v\}$, and m channels. Define $sw(A)$ to be the social welfare of allocation A. If $sw(T_v) > sw(S)$ we let $S = T_v$ be our new allocation and go back to the beginning of the while loop. Otherwise we let the payment of v be $P_v = sw(T_v) - sw(S) + b(v)$, where $b(v)$ is the sum of the valuations of l_v for the channels it received. When Algorithm 3 has passed through all links $v \in S$ it returns the best allocation found S and the payment vector P. If the algorithm finds an allocation with higher social welfare during payment calculation it restarts with this improved allocation. This can happen at most n times, since payments are calculated for at most n links.

Algorithm 3: Payment computation for GREEDYWEIGHT.

```
Let S be the allocation found by GREEDYWEIGHT
Let P_{n×1} = 0 be the payment vector
for v ∈ S do
    Let T_v be the allocation returned by GREEDYWEIGHT for links V \ {v},
    and m channels
    P_v = max(0, sw(T_v) − sw(S) + b_v)
end
```

Output:
- Allocation S.
- Payment vector P.

We know that in the final allocation S, that for any link v, $sw(T_v) \leq sw(S)$, otherwise T_v would have been chosen as an allocation instead. If l_v is bidding truthfully, it will never receive negative utility, since $u_v = b_v - P_v \geq b_v - sw(T_v) + sw(S) - b_v \geq 0$, hence it is individually rational. Furthermore it has no positive transfers, since the payment of each link v is $P_v = \max(0, sw(T_v) - sw(S) + b_v)$, and must therefore be nonnegative.

However, the payments computed by Algorithm 3 do not turn GREEDYWEIGHT into a truthful mechanism. Consider a single channel auction, $m = 1$, with a set of links $V = \{l_1, l_2, l_3, l_4\}$, with valuations $b_1 = 3, b_2 = 2, b_3 = 2 - \epsilon, b_4 = 2 - \epsilon$. Let the maximal feasible sets be $S_1 = \{l_1\}, S_2 = \{l_2\}, S_3 = \{l_3, l_4\}$ and no others. Algorithm 2 orders the links in decreasing order of valuations l_1, l_2, l_3, l_4 and allocates the channel to S_1. Algorithm 3 would then order the links $V \setminus \{l_1\}$ in decreasing order of valuations l_2, l_3, l_4 and greedily choose S_2 as the runner up. The payment of l_1 would then be $P_1 = sw(S_2) - sw(S_1) + b_1 = 2 - 3 + 3 = 2$ and utility $u_1 = 3 - 2 = 1$. The utility of other links is $u_2 = u_3 = u_4 = 0$.

Now see what happens when l_3 makes a bid $b_3^* = 2 + \epsilon$ which is higher than its true valuation $b_3 = 2 - \epsilon$. Algorithm 2 would order the links in decreasing order of valuations l_1, l_3^*, l_2, l_4 and still allocate the channel to S_1. Algorithm 3 would then order the links in $V \setminus \{l_1\}$ in decreasing order of valuations l_3^*, l_2, l_4 and finds the allocation S_3 which is better than S_1 and makes it the new allocation. Then it calculates the payments $P_3^* = sw(S_1) - sw(S_3) + b_3^* = 3 - 4 + 2 + \epsilon = 1 + \epsilon$ and $P_4 = sw(S_1) - sw(S_3) + b_4 = 3 - 4 + 2 - \epsilon = 1 - \epsilon$. Now we can see that the utility of l_3 is $u_3^* = b_v - p_3^* = 2 - \epsilon - 1 - \epsilon = 1 - 2\epsilon > u_3$, when it reports its valuation as b_3^* instead of its true valuation b_3. Thereby we have shown that Algorithm 2 with payments computed by Algorithm 3 is not a truthful mechanism.

3.3 Algorithm BUCKETLENGTH

Algorithm 4 is called BUCKETLENGTH. Initially links are divided into buckets $T_1, \ldots, T_r$ of roughly equal length, meaning that in any bucket T_t a link $u \in T_t$ is at most twice the size of any other link $v \in T_t$. Then we initialize empty sets $S_1, \ldots, S_m$ for each channel to be allocated. For each channel k, we find an allocation as follows. We initialize a temporary allocation U_t for each bucket t, which gets the allocation found by GREEDYWEIGHT with the links in the bucket as input. The valuation of link $l_w \in U_t$ for an additional channel is denoted as b_w. Then S_k is allocated to U_t with the highest social welfare.

Algorithm 4: Algorithm BUCKETLENGTH for downward sloping symmetric valuations.

```
Divide links of similar lengths into buckets T_1, ..., T_r
S_1, ..., S_m ← ∅
for k ∈ {1, ..., m} do
    for t ∈ {1, ..., r} do
        U_t ← GREEDYWEIGHT(T_t, 1)
    end
    t_0 ← arg max_{t∈{1,...,r}} (∑_{l_w∈U_t} b_w)
    S_k ← U_{t_0}
end
```

Output:

- Allocate each channel i to set S_i of links, $i = 1, \ldots, m$.

With algorithm BUCKETLENGTH we eliminate the chance of a long link v ruling out a feasible set S of smaller links of arbitrary size. Unless the valuation of v is higher than the sum of all links in S, in which case it is better to allocate to v.

3.4 Algorithm BUCKETWEIGHT

Algorithm 5 begins by initializing empty sets $S_1, \ldots, S_m$ for each channel to be allocated. Then for each channel it puts the links into buckets $T_1, \ldots, T_r$, such that for any links $u, v \in T_i$, the valuation of $b_u(1)$ of u is at most twice the valuation $b_v(1)$ of v for an additional channel. Then GREEDYLENGTH takes links in each bucket and their valuation for an additional channel and finds a temporary allocation assigned to U_t. The valuation of link $l_w \in U_t$ for an additional channel is denoted as b_w. Finally S_k gets the allocation U_i which has the highest social welfare.

Algorithm 5: Algorithm BUCKETWEIGHT for downward sloping symmetric valuations.

$S_1, \ldots, S_m \leftarrow \emptyset$
for $k \in \{1, \ldots, m\}$ **do**
  Divide links of similar valuations for an additional channel into buckets $T_1, \ldots, T_r$
  for $t \in \{1, \ldots, r\}$ **do**
    $U_t \leftarrow$ GREEDYLENGTH$(T_t, 1)$
  end
  $t_0 \leftarrow \arg\max_{t \in \{1,\ldots,r\}} (\sum_{l_w \in U_t} b_w)$
  $S_k \leftarrow U_{t_0}$
end

Output:
- Allocate each channel i to set S_i of links, $i = 1, \ldots, m$.

The problem with GREEDYLENGTH was that it finds an allocation based only on the length of the links, with no regard to their valuation. When auctioning multiple channels it allocates all the channels to the same links. The algorithm BUCKETWEIGHT applies GREEDYLENGTH on buckets of links with similar valuations, which are recomputed for each channel and hopefully results in higher social welfare than GREEDYLENGTH.

3.5 Algorithm LOCALRATIO

Algorithm 6, LOCALRATIO, is based upon a non-standard graph parameter ρ called inductive independence number.

The algorithm orders the links in the inductive independent ordering, which happens to be increasing order of length for the SINR model. Then it initializes empty sets $S_1, \ldots, S_m$ one for each channel to be allocated. For each channel k, we initialize a vector $W_{n\times 1} = 0$. Then for each link v in increasing order of length assign $W_v = \max(0, b_v(1) - \sum_{u \in Prev(v)} W_u \cdot C_{vu})$, where $Prev(v)$ is the backward neighbourhood of v. Then we go through each link v in decreasing order of length and if $W_v > 0$ and feasibility conditions hold $\forall u \in S_k \cup v$ then v is added to S_k. The channels are then allocated to the sets $S_1, \ldots, S_m$.

Algorithm 6: Algorithm LOCALRATIO for downward sloping symmetric valuations

Order links in increasing order of length, $d^*_{11} \leq d^*_{22} \leq \cdots \leq d^*_{nn}$
$S_1, \ldots, S_m \leftarrow \emptyset$
for $k \in \{1, \ldots, m\}$ **do**
  $W_{n\times 1} = 0$
  for $v = 1$ *to* n **do**
    $W_v = \max(0, b^*_v(1) - \sum_{u \in \mathit{Prev}(v)} W_u \cdot C_{vu})$
  end
  for $v = n$ *to* 1 **do**
    if $W_v > 0$ **then**
      if $\frac{P_{uu}}{\sum_{w \in S_k \cup \{l^*_v\} \setminus \{l_u\}} P_{wu} + N} \geq \beta, \forall u \in S_k \cup \{l^*_v\}$ **then**
        $S_k \leftarrow S_k \cup \{l^*_v\}$
      end
    end
  end
end

Output:
- Allocate each channel i to set S_i of links, $i = 1, \ldots, m$.

The main strength of LOCALRATIO is its approximation guarantee. It can be shown that LOCALRATIO gives a ρ-approximation to the weighted capacity problem in graph-based models (Ye & Borodin, 2009), and ρ for conflict graphs in the SINR model are in the order of $\mathcal{O}(\log(n))$ (Hoefer & Kesselheim, 2012).

3.5.1 Payment Computation for LOCALRATIO

The payment computation for LOCALRATIO is also inspired by the VCG mechanism, which is truthful but relies on finding the optimal allocation (Vickrey, 1961; Clarke, 1971; Groves, 1973). We have no guarantee that we find the optimal allocation with LOCALRATIO although it may happen once in a while.

The payment computation for LOCALRATIO goes as follows. Let S be the allocation found by LOCALRATIO for links V and m channels. Let $P_{n\times 1} = 0$ be a vector used to store the payment of each link. For all $v \in S$ in any order, let T_v be the allocation found by LOCALRATIO for links $V \setminus \{v\}$, and m channels. Define $sw(A)$ to be the social welfare of allocation A. If $sw(T_v) > sw(S)$, let $S = T_v$ be our new allocation and go back to the beginning of the while loop. Otherwise we let the payment of v be $P_v = sw(T_v) - sw(S) + b(v)$, where $b(v)$ is the sum of the valuations of v for the channels it received. When Algorithm 3 has passed through all links $v \in S$, it returns the best allocation S found and the payment vector P. If the algorithm finds an allocation with higher social welfare during payment calculation it restarts with this improved allocation. This can happen at most n times, since payments are calculated for at most n links.

Algorithm 7: Payment computation for LOCALRATIO.

```
Let S be the allocation found by LOCALRATIO
Let P_{n×1} = 0 be the payment vector
for v ∈ S do
    Let T_v be the allocation returned by LOCALRATIO for
    links V \ {v}, and m channels
    P_v = max(0, sw(T_v) − sw(S) + b_v)
end
```

Output:

- Allocation S.
- Payment vector P.

We know that in the final allocation S, that for any link v, that $sw(T_v) \leq sw(S)$, otherwise T_v would have been chosen as an allocation instead. If v is bidding truthfully, than it will never receive negative utility, since $u_v = b_v - P_v \geq b_v - sw(T_v) + sw(S) - b_v \geq 0$, hence it is individually rational. Furthermore it has no positive transfers, since the payment of each link v is $\max(0, sw(T_v) - sw(S) + b_v)$, and must therefore be nonnegative.

The payments computed by Algorithm 7 do not however turn LOCALRATIO into a truthful mechanism. Consider a single channel auction, $m = 1$, with a set of links $V = \{l_1, l_2\}$, where l_1 is the shorter link and conflict caused from l_1 on l_2, and vice versa is $C_{1,2} = C_{2,1} = \epsilon$. Let their true valuations be $b_1 = 1$ and $b_2 = \epsilon$, respectively.

Algorithm 6 excludes l_2 from consideration in the forward pass of the algorithm and selects l_1 into the allocation S in the backward pass. The social welfare of S is $sw(S) = 1$, and Algorithm 7 calculates the payment of l_1 as $P_1 = \epsilon$. Their utilities are then, $u_1 = 1 - \epsilon$ and $u_2 = 0$, respectively.

Let's assume that l_2 makes a bid $b^* = 2\epsilon$, which is higher than its true valuation $b_v = \epsilon$. Now Algorithm 6 does not exclude it in the forward pass, and both links are accepted into the set S^* in the backward pass. The social welfare of S^* in this case is $sw(S^*) = 1 + \epsilon$ and the payment calculated by Algorithm 7 are $P_1^* = \epsilon - 1 - \epsilon + 1 = 0$ and $P_2^* = 1 - 1 - \epsilon + \epsilon = 0$, respectively. Their utilities are then, $u_1^* = 1$ and $u_2^* = \epsilon$, respectively.

Thus we have shown that l_2 benefits from lying, since $u_2^* > u_2$ and Algorithm 7 does not turn Algorithm 6 into a truthful mechanism.

3.6 Algorithm GREEDYINTERFERENCE

Algorithm 8 is called GREEDYINTERFERENCE and is based on the intuition that it is generally good to accept robust links with high valuation. Let L be the set of bidding links and $l_v \in L$. Then l_v gets a value ω_v, which is its SINR value scaled with its valuation for an additional channel and the valuations of other links $l_w \in L \setminus \{l_v\}$.

Definition 5. *Assign to each link l_v a value ω_v, called weighted SINR. It is the SINR value of l_v if all links $l_w \in L \setminus \{l_v\}$ would transmit scaled with the valuation of each link,*

$$\omega_v = \frac{P_{vv} \cdot b_v}{\sum_{w \in L \setminus \{v\}} P_{wv} \cdot b_w}, \tag{3.1}$$

where L is the set of all bidding links and b_v is valuation of l_v for an additional channel.

The algorithm begins by initializing empty sets $S_1, \ldots, S_m$, one for each channel to be allocated. For each channel k, the algorithm greedily adds links in decreasing order of weighted SINR ω if feasibility conditions hold. Finally the algorithm allocates one channel to each set $S_1, \ldots, S_m$.

Algorithm 8: Algorithm GREEDYINTERFERENCE for downward sloping symmetric valuations.

$S_1, \ldots, S_m \leftarrow \emptyset$
for $k \in \{1, \ldots, m\}$ **do**
  Order links in decreasing order of weighted SINR,
  $\omega_1^* \geq \omega_2^* \geq \cdots \geq \omega_n^*$
  for $v = 1$ *to* n **do**
    if $\frac{P_{uu}}{\sum_{w \in S_k \cup \{l_v^*\} \setminus \{l_u\}} P_{wu} + N} \geq \beta, \forall u \in S_k \cup \{l_v^*\}$ **then**
      $S_k \leftarrow S_k \cup \{l_v^*\}$
    end
  end
end

Output:
- Allocate each channel i to set S_i of links, $i = 1, \ldots, m$.

Although GREEDYINTERFERENCE is meant to overcome the arbitrarily poor performance of GREEDYLENGTH and GREEDYWEIGHT in special cases, it has no proven approximation guarantee.

3.7 Linear Programming Algorithm (LP)

Next algorithm is called LP and consists of two phases. First it needs to solve a Linear Program (LP) which takes inductive independence number ρ as parameter. Its objective function is to maximize the sum of valuations for all links for the number of channels they are allocated, fractional allocations allowed. It is subject to constraints that the sum of conflict caused by links u in the backward neighbourhood of any link v multiplied by the number of channels allocated is less than or equal to $\rho \cdot m$. The different fractional allocations of number of channels to each link may not exceed 1 and allocations must be nonnegative.

The LP relaxation reads

$$\begin{array}{llr}
\text{Maximize} & \displaystyle\sum_{v \in V} \sum_{i=1}^{m} b_v(i) \cdot x_{v,i} & \\
\text{subject to} & \displaystyle\sum_{\substack{u \in V \\ \pi(u) < \pi(v)}} \sum_{i=1}^{m} i \cdot C(u,v) \cdot x_{u,i} \le \rho \cdot m & \forall v \in V \\
& \displaystyle\sum_{i=1}^{m} x_{v,i} \le 1 & \forall v \in V \\
& x_{v,i} \ge 0 & \forall v \in V, i \in [m],
\end{array}$$

where $b_v(i)$ is the valuation of v for its i-th channel, x is the fractional allocation, C is the conflict graph, ρ the inductive independence number, m the number of channels and V the set of bidding links.

The first phase of the algorithm is formulated in the same way as suggested by the authors (Hoefer & Kesselheim, 2012). Note that this relaxation does not describe the channel allocation problem exactly, an integral solution might not be feasible for the channel allocation problem.

We then use a different rounding phase than proposed, which seems to return allocations with higher social welfare in simulations. We use Algorithm 9 to round the fractional LP solution. First initialize empty sets $S_1, \dots, S_m$. Then we run the following randomized algorithm a thousand times and return the allocation with the highest social welfare. Initialize empty temporary sets $T_1, \dots, T_m$. Take the solution x, returned by the LP phase and round each value $x_{v,i}$ to 1 with probability $x_{v,i}$ and 0 otherwise. Initialize a vector R and let it have i copies of v for each $x_{v,i}$ equal to 1 after the rounding phase. Then for all $v \in R$ in random order check if it can be added to any channel without violating SINR feasibility constraints and add it if possible. When all $v \in R$ have been considered compare the social welfare of the temporary solution T with S and if it has higher social welfare let T be the new allocation S. After running this a thousand times the alogrithm allocates each channel $1, \dots, m$ to a correspondings set of links $S_1, \dots, S_m$.

Algorithm 9: Rounding LP solution for downward sloping symmetric valuations.

$S_1, \dots, S_m \leftarrow \emptyset$
for $t \in \{1, \dots, 1000\}$ **do**
  $T_1, \dots, T_m \leftarrow \emptyset$
  $x_{v,i} = 1$ with probability $x_{v,i}$, 0 otherwise
  Let R have i copies of v for each $x_{v,i}$, where $x_{v,i} = 1$.
  for $v \in R$ *in random order* **do**
    for $k \in \{1, \dots m\}$ *in random order* **do**
      if $\frac{P_{uu}}{\sum_{w \in T_k \cup \{l_v^*\} \setminus \{l_u\}} P_{wu} + N} \geq \beta, \forall u \in T_k \cup \{l_v^*\}$ **then**
        $T_k \leftarrow T_k \cup \{l_v^*\}$
        break
      end
    end
  end
  if $sw(T) > sw(S)$ **then**
    $S \leftarrow T$
  end
end

Output:
- Allocate each channel i to set S_i of links, $i = 1, \dots, m$.

The main strength of the LP algorithm with the original rounding phase is that it produces social welfare of at least $\frac{b^*}{16\sqrt{m\rho}\lceil \log n \rceil}$ in expectation (Hoefer et al., 2011). By using our rounding phase instead we achieve higher social welfare in simulations although we can offer no approximation expectation or guarantee.

3.8 Optimal Solution

We formulated a mixed integer program (MIP) which finds an optimal solution to the channel allocation problem, if allowed enough time.

Let V be the set of links, I be the set of channels and let $b_v(i)$ be the valuation of link v for its i-th channel. The binary variable $y_{v,i}$ takes the value 1 if l_v gets its i-th channel, 0 otherwise. The binary variable $x_{v,i}$ takes the value 1 if v gets channel i, 0 otherwise. Let M be a large number, C be the conflict graph and β represent the minimum SINR for a feasible set.

The MIP reads

$$\begin{array}{lrl}
\text{Maximize} & & \displaystyle\sum_{v\in V}\sum_{i\in I} b_v(i)\cdot y_{v,i} \\
\text{subject to} & \displaystyle M\cdot x_{v,i} + \sum_{u\in V\setminus\{v\}} C(u,v)\cdot x_{u,i} + \frac{N}{P_{vv}} \le M + \frac{1}{\beta} & \forall v\in V, i\in I \\
& \displaystyle\sum_{i\in I} y_{v,i} - \sum_{i\in I} x_{v,i} = 0 & \forall v\in V \\
& x_{v,i}, y_{v,i} \in \{0,1\} & \forall v\in V, i\in I.
\end{array}$$

The objective function is to maximize the sum of valuations over all links and is subject to the constraints of SINR feasibility, formulated by using the conflict graph C. It is also subject to the constraint that x and y are binary matrices and that for all v we have $\sum_{i\in I} y_{v,i} - \sum_{i\in I} x_{v,i} = 0$.

The obvious advantage of the optimal MIP is that it returns an allocation with maximum social welfare. If the links are few and there is enough time, this algorithm should definitely be applied with VCG payments. However for a large number of links this algorithm is impractical and previously mentioned approximation algorithms are more suitable.

Chapter 4

Simulations

We generate the topology for our simulations by using a model inspired by the preferential attachment model from network theory (Barabási & Albert, 1999) and further developed for wireless communication in the SINR model (Ásgeirsson et al., 2012). To match phenomena seen in real-world complex networks we generate graphs with power-law degree distributions. For each simulation we generated a random instance of links as follows. We begin with an empty $x \times y$ plane. For the first link l_1 we find a random location in the plane for the sender s_1, a random angle $\theta = [0, 2\pi]$ to send to and place the receiver r_1 at a random distance between $d_{\min}$ and $d_{\max}$. For all other links l_v, with probability p we place the center of l_v close to some other previously placed link, distance randomly generated from a Pareto distribution with minimum distance 0.1. Then choose a random angle $\theta = [0, 2\pi]$ and distance d_{vv} between $d_{\min}$ and $d_{\max}$. Then we place the sender s_v at distance $\frac{d_{vv}}{2}$ in direction θ from the center and the receiver r_v is placed at distance $\frac{d_{vv}}{2}$ in the opposite direction from the center. With probability $1 - p$, l_v is placed randomly in the plane like the first link. Each link l_v transmits with mean power $P_v = d_{vv}^{\alpha/2}$. We note that $p = 0$ gives uniform random distribution.

Each link was then given downward sloping and symmetric valuations. This means that they only care about how many channels they are allocated and they are willing to pay less for each additional channel. Furthermore they do not care which of the other links are allocated channels. The valuation of each link l_v is generated in the following way. Let r_v be a random number between 0 and 1, and multiply it with the transmission power P_v of l_v. Then let $b_{v,k} = r_v + r_v \cdot \sin(\frac{3\pi}{2}) \cdot \frac{k-1}{m}$ be the valuation of l_v for its k-th additional channel in an m channel auction.

In our simulations we place the links on a 20×20 plane ($x = 20$ and $y = 20$). The minimum length of a link $d_{\min} = 0.01$ and maximum length $d_{\max} = 5.12$ was chosen such that there is a large range in the length of links but the number of buckets used in BUCKETLENGTH is limited to ten. We chose the path loss coefficient $\alpha = 2.1$, minimum SINR $\beta = 1$ and ambient noise $N = 0$ (see Table 4.1).

Length of the plane (x)	20
Width of the plane (y)	20
Minimum length of link ($d_{\min}$)	0.01
Maximum length of link ($d_{\max}$)	5.12
Path loss coefficient (α)	2.1
Minimum SINR (β)	1
Ambient noise (N)	0

Table 4.1: Constants used in all simulations.

In the first part of the simulations we compared social welfare of the fractional solution of the LP algorithm to the rounded solution and to an optimal solution. The aim was to find a value for the inductive independence number ρ which is an input in the LP algorithm that gives the highest social welfare of the rounded LP solution. In the second part of the simulations we compare the approximation algorithms when we vary the number of links, in the third part we vary the number of channels and in the fourth part we vary the probability of attaching. The aim of the second, third and fourth part of the simulations is to give an auctioneer the opportunity to compare the algorithms performance in different situations, such that he can choose the algorithm that fits best with his current situation. In the fifth part of the simulations we compare the social welfare obtained by GREEDYWEIGHT and LOCALRATIO after payment computation. We also compare the total payments made when using these algorithms, which may be of interest to some auctioneers. In the sixth part of the simulations we compare the social welfare obtained by GREEDYWEIGHT and LOCALRATIO for varying valuation functions. The social welfare of an optimal allocation is presented as comparison where it is computable in reasonable time.

4.1 Finding a Good ρ

One of the inputs to the LP algorithm is the inductive independence number ρ. Since we do not know the inductive independence number of the graph we run simulations and examine the effects of changing the value ρ. We run the simulations with $n = 40$ links, $m = 4$ channels and probability of attaching $p = 0.4$. For each value $\rho = \{1, 2, \ldots, 10\}$ we repeat the simulations $k = 80$ times (see Table 4.2).

Number of links (n)	40
Number of channels (m)	4
Probability of attaching (p)	0.4
Number of trials (k)	80

Table 4.2: Constants used in the simulation for varying ρ.

The optimal allocation had mean social welfare $\mu = 59.0$ with standard deviation $\sigma = 7.5$. The social welfare of the optimal solution is not normally distributed according to Kolmogorov-Smirnov test. Table 4.3 shows some characteristics of the social welfare of the optimal solution in our $k = 80$ trials. We can see that the optimal social welfare is in the range from 44.9 to 73.5.

μ	σ	Normally distributed	Min	Q_1	Median	Q_3	Max
59.0	7.5	No	44.9	53.2	58.7	65.2	73.5

Table 4.3: Social welfare of the optimal solution.

If the social welfare of the LP fractional solution is lower than the social welfare of the optimal solution some links in the optimal solution never get a chance in the rounding phase. On the other hand if the social welfare of the fractional solution is too much above the social welfare of the optimal solution we are including too many links which makes it unlikely that we find the optimal solution in the randomized rounding phase. Looking at the extremes gives us deeper understanding of how the algorithm works. If we let $\rho = 0$ than no link is allocated any channel in the LP fractional solution, making the randomized rounding phase useless. However if we let $\rho = \infty$ then all links are allocated all channels and it's only up to the rounding phase to allocate the channels, making the LP phase useless.

In Figure 4.1 we can see that for $\rho = 1$ the social welfare of the fractional LP solution is on average lower than the social welfare of the optimal solution. For all other values tested for ρ the social welfare of the LP fractional solution is well above the social welfare of the optimal solution on average. However we see that increasing the value of ρ only pays off up to $\rho = 4$, then the average social welfare of the rounded solution declines with increasing values of ρ. The difference between social welfare obtained for different values of ρ is small and their is lot of variance, we can therefore not exclude the possibility that the differences observed are due to random variance (see Table A.1).

The finding of $\rho = 4$ giving the best results is in line with previous results which state that $\rho = \mathcal{O}(\log(n))$ for wireless links in the SINR model (Hoefer & Kesselheim, 2012), since $\log(40) \approx 3.69$. Therefore we let $\rho = \log(n)$ in other parts of the simulations.

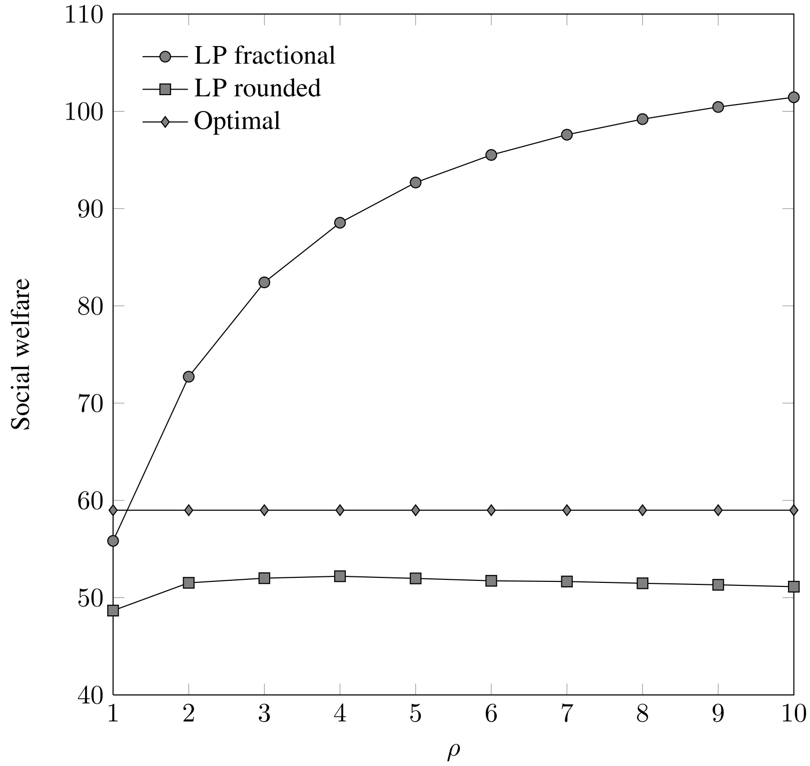

Figure 4.1: Social welfare of the LP fractional, LP rounded and optimal solution with varying values for ρ.

4.2 Number of Links

What will probably vary the most is the number of links n. Therefore it is very important for the auctioneer to know how the performance of the approximation algorithms varies with number of links. For this part of the simulation we keep the number of channels $m = 4$ constant and the probability of attaching $p = 0.4$. We run the algorithms on $k = 80$ randomly generated instances for each number of links $n = \{10, 20, \ldots, 100\}$ tested (see Table 4.4).

Number of channels (m)	4
Probability of attaching (p)	0.4
Number of trials (k)	80

Table 4.4: Constants used in the simulation for varying number of links (n).

In Figure 4.2 we can see the mean social welfare of the approximation algorithms. We see that GREEDYLENGTH returns the lowest social welfare for all number of links tested, this should be expected and is partly due to how the valuations for each link are generated (see Table A.2).

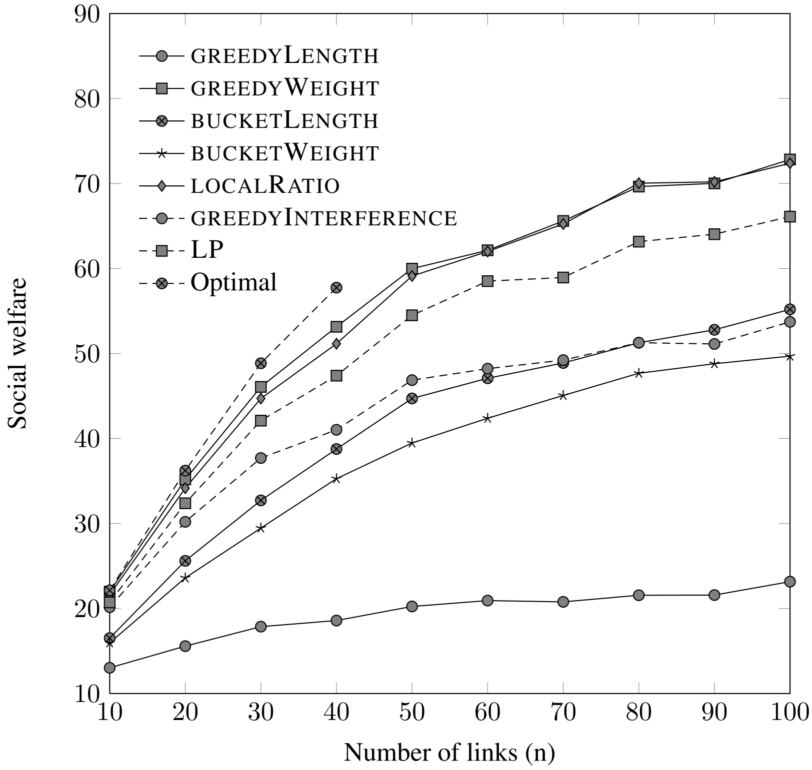

Figure 4.2: Social welfare of the algorithms when varying number of links.

The second worst is bucketWeight (see Table A.5), which runs greedyLength on buckets of links with similar weights and picks the best one. It is interesting however to see how much bucketWeight improves on greedyLength.

Ranking fourth and fifth in average social welfare are bucketLength and greedyInterference, with greedyInterference producing higher social welfare on average for $n = \{10, 20, ..., 60\}$, equally high for $n = 70$ and lower for $n = \{90, 100\}$. Possibly greedyInterference is better for small number of links, but then bucketLength is better, with more links in each bucket. Further simulations with higher number of links could support or refute that hypothesis, but the variance is too large to be certain (see Tables A.4 and A.7).

The LP algorithm ranks third, which relies on the inductive independence number ρ selected and heavily relies on the rounding phase. By choosing the right ρ and improving the rounding phase there might be room for improvement and higher ranking for the LP (see Table A.8).

Ranking first and second are greedyWeight and localRatio, with greedyWeight producing higher social welfare on average for all values of n except $n = \{80, 90\}$. One may be tempted to state that GREEDYWEIGHT is the best but the variance is large and without further simulations it remains uncertain (see Tables A.3 and A.6).

The results of these simulations imply that the auctioneer should choose either GREEDYWEIGHT or LOCALRATIO no matter the number of links. The difference between the two algorithms is small, with their relative ranking interchanging, but always higher than all other algorithms. The difference between the social welfare obtained by GREEDYWEIGHT and LOCALRATIO is small, but seems to be increasing with the number of channels (see Table 4.3 for social welfare of the optimal solution).

4.3 Number of Channels

The number of channels m the auctioneer can put up for auction may vary. Being able to compare the social welfare of allocations output by the approximation algorithms with different number of channels is valuable information to the auctioneer who wishes to get as high social welfare as possible. In this part of the simulations we keep the number of links $n = 40$ constant as well as the probability of attaching $p = 0.4$. We run the algorithms on $k = 80$ randomly generated instances for each number of channels $m = \{1, 2, \ldots, 10\}$ (see Table 4.5).

Number of links (n)	40
Probability of attaching (p)	0.4
Number of trials (k)	80

Table 4.5: Constants used in the simulation for varying number of channels.

In Figure 4.3 we see that the ranking of the algorithms is very similar as when we varied the number of links n. The worst social welfare of the approximation algorithms is GREEDYLENGTH and second worst BUCKETWEIGHT (see Tables A.10 and A.13).

Consistently ranking fourth and fifth are GREEDYINTERFERENCE and BUCKETLENGTH, respectively. This differs from what we saw when we varied the number of links n, where GREEDYINTERFERENCE produced higher social welfare for small n but BUCKETLENGTH was better for larger n (see Tables A.15 and A.12).

As previously LP ranks third (see Table A.16). The difference between GREEDYWEIGHT and LOCALRATIO is small as before, but now GREEDYWEIGHT consistently produces higher social welfare than LOCALRATIO. It is therefore safe to rank GREEDYWEIGHT first and LOCALRATIO second in this part of the simulations (see Tables A.14 and A.11).

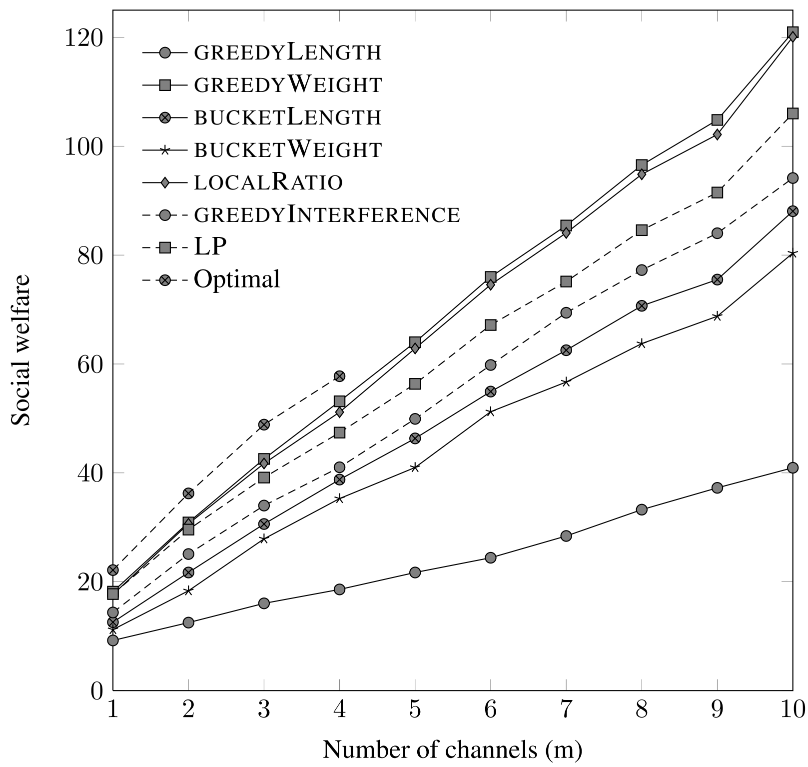

Figure 4.3: Social welfare of the algorithms when varying number of channels.

It is interesting to see that the change in social welfare as we increase the number of channels is close to linear for all algorithms. What differs is their starting point and the slope of the lines. This is probably a product of how the valuation functions are computed. With different valuation function we would see a completely different graph.

In this part of the simulations GREEDYWEIGHT consistently ranks first, with LOCALRATIO close behind. These results give the auctioneer no reason to consider other algorithms if his goal is to maximize social welfare. Furthermore we can see that the social welfare of the optimal solution, for $n \leq 40$ is not far away (see Table A.17).

4.4 Probability of Attaching

Ranging from $p = 0$ with all links randomly distributed on the plane, to $p = 1$ where all links attach to previous links, the probability of attaching may be varied to model different kinds of situations. The values between 0 and 1 produce more interesting situations, going from rural to urban as p approaches 1. For this part of the simulations we keep the number of links $n = 40$ constant as well as the number of channels $m = 4$. We run the algorithms on $k = 80$ randomly generated instances for probability of attaching $p = \{0.1, 0.2, \ldots, 1.0\}$ (see Table 4.6).

Number of links (n)	40
Number of channels (m)	4
Number of trials (k)	80

Table 4.6: Constants used in the simulation for varying probability of attaching (p).

In Figure 4.4 we see the effects of changing p on the social welfare produced by our approximation algorithms. For values of $p \leq 0.4$ the effect is insubstantial, but as p gets closer to 1 the social welfare produced drops quickly.

We see that GREEDYLENGTH produces the lowest social welfare for all values of p. Ranking fifth and sixth are BUCKETLENGTH and BUCKETWEIGHT, respectively, except when p is close to 1.0.

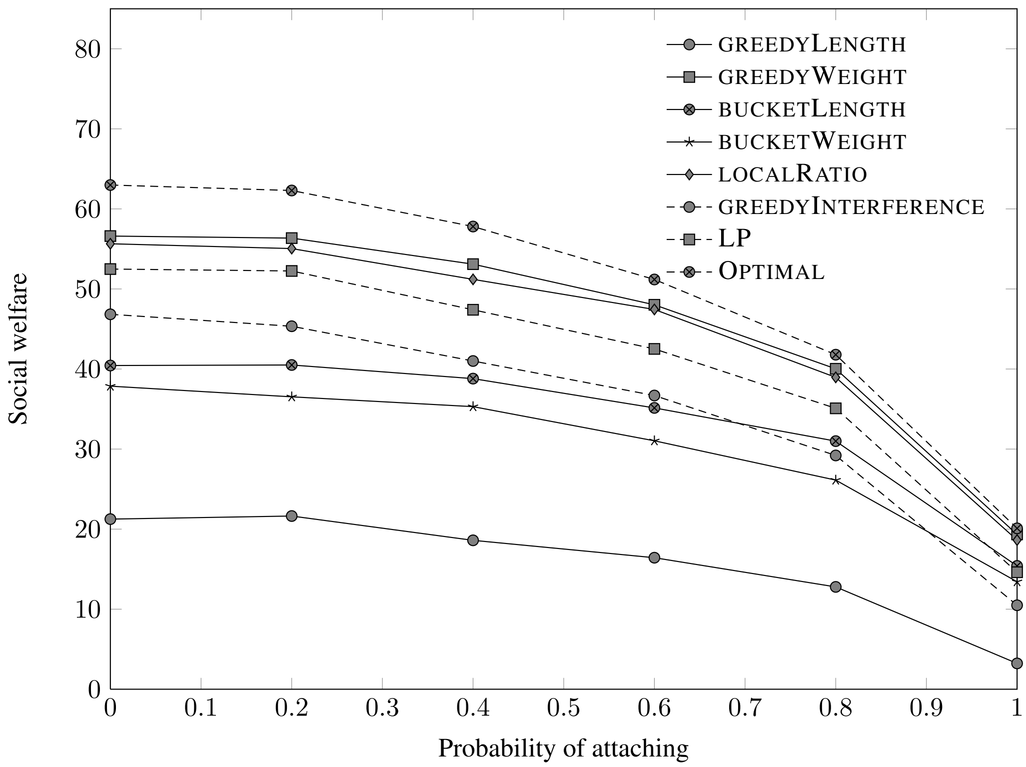

Figure 4.4: Social welfare of the algorithms when varying the probability of attaching.

We see that when p gets closer to 1.0 the social welfare of GREEDYINTERFERENCE and LP drops much faster than the social welfare of the other approximation algorithms, with GREEDYINTERFERENCE ranking sixth and LP ranking fourth when $p = 1.0$. For lower values of p LP and GREEDYINTERFERENCE rank third and fourth respectively. Consistently ranking first and second, without interchanging are GREEDYWEIGHT and LOCALRATIO, respectively (see Tables A.18, A.19, A.20, A.21, A.22, A.23 and A.24).

We still see GREEDYWEIGHT and LOCALRATIO producing the highest social welfare, for all values of p giving the decision maker little reason to consider any of the other approximation algorithms presented. Furthermore we see that the social welfare of GREEDYWEIGHT and LOCALRATIO is close to the optimal social welfare especially for p close to 1.0 (see Table A.25).

4.5 Payments for GREEDYWEIGHT and LOCALRATIO

We tried varying the number of links, channels and probability of attaching. The ranking of the average social welfare produced by the algorithms stayed roughly the same as we explored these different dimensions with GREEDYWEIGHT and LOCALRATIO ranking first and second, sometimes interchanging but always outperforming the other approximation algorithms. Our recommendations to the auctioneer are therefore to use one of these two no matter the number of links, channels or the probability of attaching.

We vary the number of links $n = \{10, 20, \ldots, 100\}$, but keep the number of channels constant $m = 4$ and the probability of attaching $p = 0.4$. We randomly generate $k = 80$ instances to test the algorithms (see Table 4.7).

Number of channels (m)	4
Probability of attaching (p)	0.4
Number of trials (k)	80

Table 4.7: Constants used in the simulation for payment computation of GREEDYWEIGHT and LOCALRATIO.

In Figure 4.5 we see the average social welfare produced by the two algorithms as well as the average total payment computed. We see that the social welfare of GREEDYWEIGHT is on average higher for smaller n but the social welfare of LOCALRATIO is on average higher for larger n although the difference is small. The total payment computed is on average higher for GREEDYWEIGHT for all number of links n. It is interesting to see that the average payment rises up to $n = 40$, but then it declines. It would be interesting to see what happens if we increase the number of links even further, whether the total payment approaches 0 (see Tables A.26 and A.28).

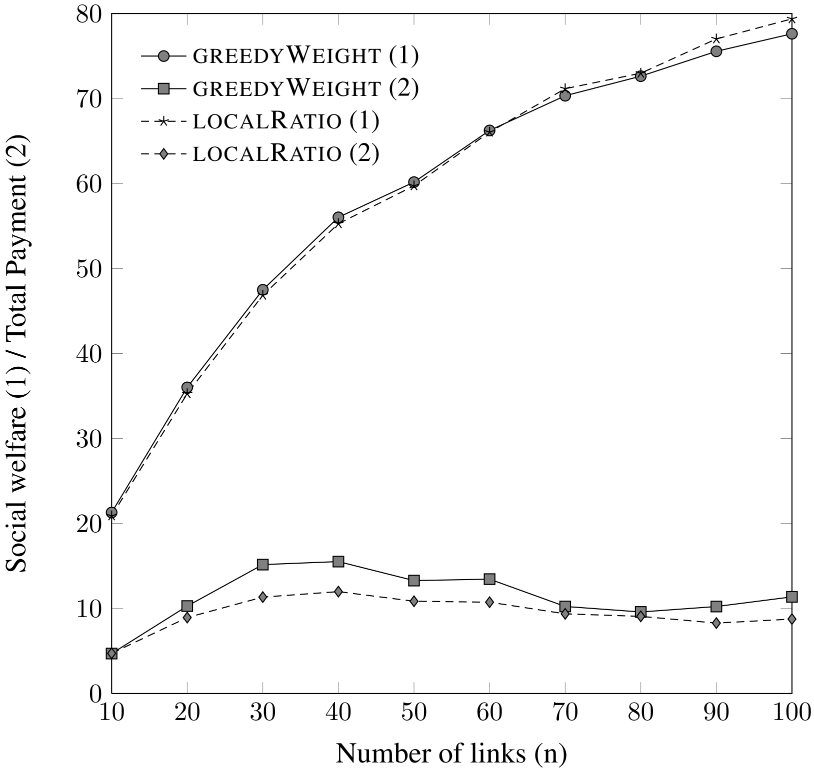

Figure 4.5: Social welfare (1) and total payment (2) of GREEDYWEIGHT and LOCALRATIO when varying number of channels.

If we find an allocation with higher social welfare than the one given by the algorithm we change our allocation to the improved one. In theory this could happen up to n times, but it seems in practice that this happens seldom. In Figure 4.6 we can see how often this happens on average for each algorithm. We see that the number of changes made on average is generally higher for LOCALRATIO than GREEDYWEIGHT. The only downside to this is increased time for payment computation. By looking at the graph it seems likely that the number of changes made while computing payment follows some logarithmic function at least in this simulation setup (see Tables A.27 and A.29).

4.6 Varying Valuations

To see the effect of different kinds of valuations, we let each link v have valuation $b_v = d_{vv}^{\lambda}$, and vary the value of $\lambda = \{0.0, 0.1, \ldots, 2.1\}$. Note that the there is no randomness in the valuation of links, it is solely based on its length.

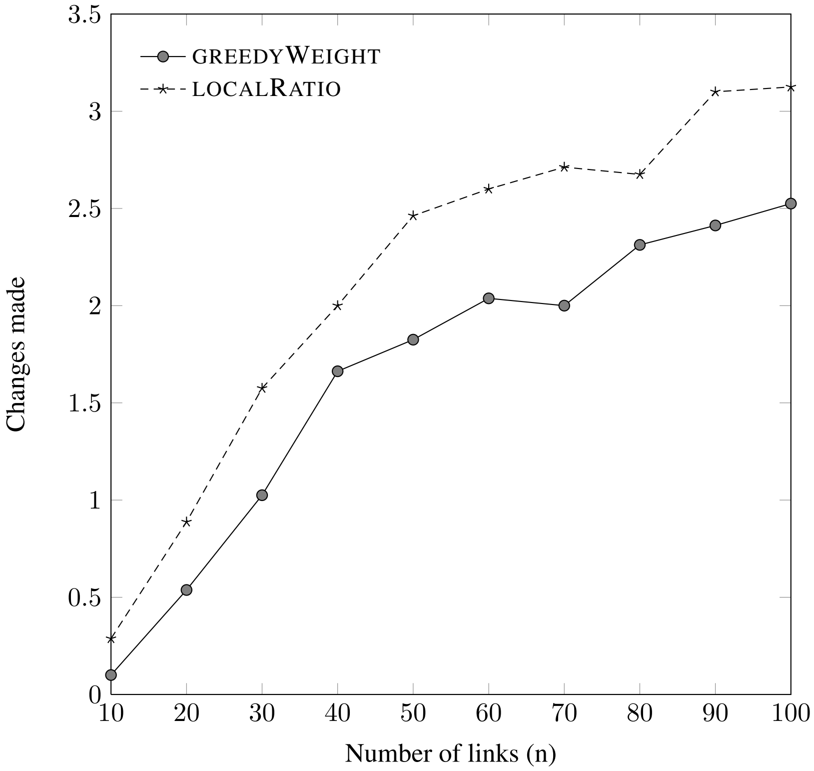

Figure 4.6: Number of changes made while calculating payment of GREEDYWEIGHT and LOCALRATIO.

Ranging from $\lambda = 0$ with all links having the same valuation, to $p = 2.1$ where valuations are the same as in other parts of the simulations. We expect lower values of λ to produce more difficulties for GREEDYWEIGHT as longer links valuations are only slightly higher than that of shorter links. Algorithm GREEDYWEIGHT will go through the links in decreasing order of length and possibly allow long links with only slightly higher valuation than short links into the allocation S.

For this part of the simulations we increase the number of links to $n = 200$ and decrease the number of channels to $m = 1$. We run the algorithms on $k = 80$ randomly generated instances for exponents $\lambda = \{0.0, 0.1, \ldots, 2.1\}$ (see Table 4.8).

Number of links (n)	200
Number of channels (m)	1
Probability of attaching (p)	0.4
Number of trials (k)	80

Table 4.8: Constants used in the simulation for varying exponent λ, used in valuations generation.

In Figure 4.7 we see that there is little difference between GREEDYWEIGHT and LOCALRATIO when $\lambda = 0.0$. This is the special case where all links have the same valuations.

Algorithm GREEDYWEIGHT goes through the link in some order not specified since there is no defined tie-breaker. Algorithm LOCALRATIO goes through the links in increasing order of length in the forward pass and decreasing order of length in the backward pass as before. In Tables A.30 and A.31 you can see the social welfare of GREEDYWEIGHT and LOCALRATIO, respectively. In all cases the social welfare is on average higher for LOCALRATIO.

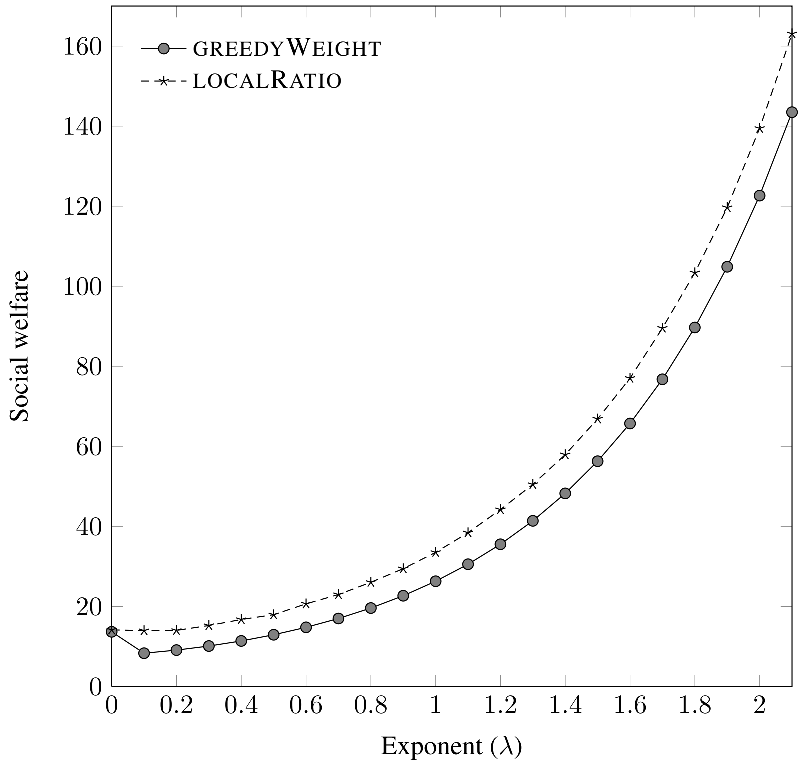

Figure 4.7: Social welfare of GREEDYWEIGHT and LOCALRATIO when varying λ to generate valuations.

This part of the simulations was meant to show that GREEDYWEIGHT will not outperform LOCALRATIO in all cases. Algorithm GREEDYWEIGHT is vulnerable to special cases where long links have slightly higher valuations. These special cases are likely to come up in these simulations since long links are given slightly higher valuations and the randomness in generation of valuations is removed. Although GREEDYWEIGHT does not perform arbitrarily badly it is considerably far off from LOCALRATIO.

Further simulations were performed to know how much effect removing randomness in valuation generation. Let B_{random} be valuations generated as described in the beginning of Chapter 4 and $B_{\text{deterministic}}$ be valuations generated in the same way except $r = 1$, instead of a random value between 0 and 1. Then we define $B_\delta = \delta \cdot B_{\text{random}} + (1 - \delta) \cdot B_{\text{deterministic}}$, where delta can take any value between 0 and 1 and determine the level of randomness in valuation generations.

In this part of the simulations we keep the number of links $n = 200$, the number of channels $m = 4$, probability of attaching $p = 0.4$ and run $k = 80$ trials for each value of δ (see Table 4.9).

Number of links (n)	200
Number of channels (m)	4
Probability of attaching (p)	0.4
Number of trials (k)	80

Table 4.9: Constants used in the simulation for varying δ used to determine the randomness in valuation generation.

In Figure 4.8 we can see that LOCALRATIO produces higher social welfare for all values of δ tested. We can also see that the difference between the algorithms is larger for low values of δ, when $B_{\text{deterministic}}$ is given relatively more weight (see Tables A.32 and A.33).

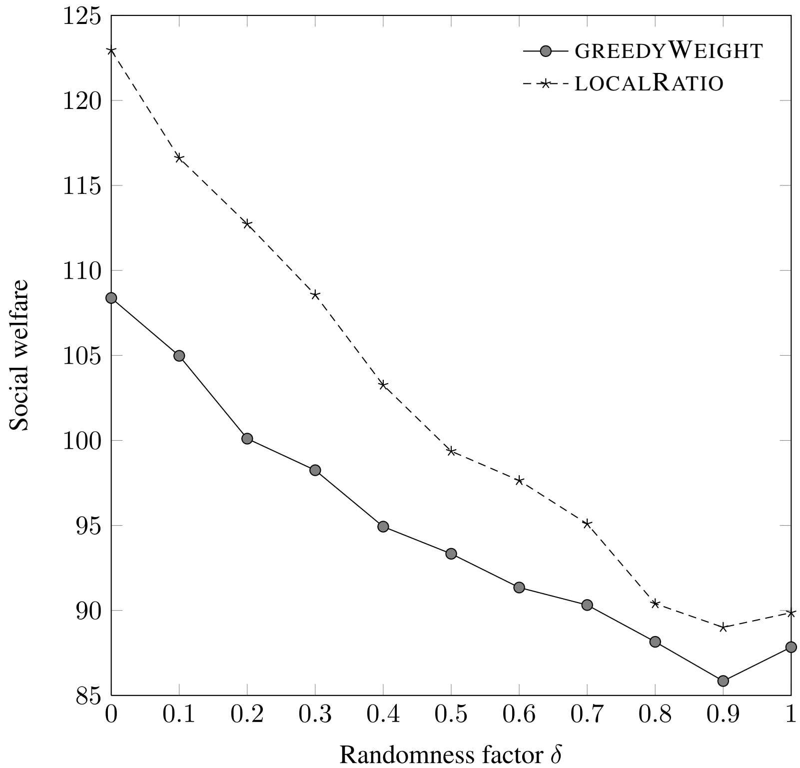

Figure 4.8: Social welfare of GREEDYWEIGHT and LOCALRATIO when varying randomness factor δ in valuation generation.

Chapter 5

Conclusions and Future Work

We have presented several algorithms that find an allocation that approximates the social welfare of a socially efficient allocation. Most of them have no proven approximation guarantee, but that does not mean that they will not do well most of the time in practice.

We explored four dimensions, through simulation, that are likely to vary in a spectrum auction, namely, number of links n, number of channels m, probability of attachment p and exponent to generate valuations λ. For all values tested for these three dimensions, two approximation algorithms consistently produced allocations with higher social welfare than the other approximation algorithms tested, that is GREEDYWEIGHT and LOCALRATIO (Algorithm 2 and 6, respectively). The social welfare obtained by these algorithms was close to the social welfare of an optimal solution for $n \leq 40$ and $m \leq 4$, but we do not know what happens for larger instances. We also see that the gap between the approximation algorithms and social welfare of an optimal solution is larger for sparse instances, that is with probability of attaching p close to 0.

Ranking third in most situations was LP (LP formulation in Section 3.7 combined with Algorithm 9). In Figure 4.1 we can see the effects of varying ρ. The effect could mainly be seen in the social welfare of the LP fractional solution. The social welfare of the rounded solution did not differ much with varying values for ρ. We can also compare the optimal social welfare to the fractional and rounded LP solution. We can see that when $\rho = 1$ the optimal social welfare is higher than the social welfare of the fractional LP solution, which tells us that the optimal solution has already been excluded.

The social welfare of the allocation found by LP was not far off from GREEDYWEIGHT and LOCALRATIO, but considering that we need to solve an LP and go through a time consuming rounding phase to find the allocation, we would want to see higher social welfare. There is however a method that turns the LP algorithm into a truthful in expectation mechanism, producing a α-approximation to the optimal social welfare, where $\alpha = 8 \cdot \sqrt{m} \cdot \rho$ (Hoefer & Kesselheim, 2012; Lavi & Swamy, 2005). The LP algorithm may also be better able to handle different kinds of valuation functions (Hoefer et al., 2011).

Producing the fourth highest social welfare in most situations was GREEDYINTERFERENCE (Algorithm 8). It is based on the intuition that it is generally good to accept links which have high valuation relative to the sum of interference it causes on all other links. In a way this algorithm is similar to LOCALRATIO. A key difference between LOCALRATIO and GREEDYINTERFERENCE is the notion of backward neighbourhood. The algorithm LOCALRATIO cleverly uses the idea of backward neighbourhood and each link considers only interference coming from those, GREEDYINTERFERENCE on the other hand is considering the interference each link causes to smaller links which may not be of importance.

Ranking fifth and sixth, are BUCKETLENGTH (Algorithm 4) and BUCKETWEIGHT (Algorithm 5), respectively. Both algorithms rely on the same principle, divide the links into buckets of a similar characteristic, find a good allocation for each bucket and return the best bucket as the final allocation. It is often a problem with bucket algorithms that many links are excluded from the beginning by falling into the wrong bucket. The algorithm BUCKETLENGTH produces significantly lower social welfare than the algorithm it uses, GREEDYWEIGHT. This is perhaps what should be expected from randomly generated instances, but this approach would overcome the special case where the social welfare of GREEDYWEIGHT is arbitrarily far off from the optimal. The other bucket algorithm, BUCKETWEIGHT produces significantly higher social welfare than the algorithm GREEDYLENGTH which it uses to find an allocation for each bucket. This is in part due to how valuations are generated.

The algorithm producing the lowest social welfare in all simulations was GREEDYLENGTH (Algorithm 1). The algorithm is initially designed to approximate the unweighted capacity and scheduling problem (Goussievskaia et al., 2009). Especially in the multi channel auction $m > 1$, it performs poorly, since it allocates all channels to the same links, but the valuations are constructed such that links are willing to pay less and less for each additional channel.

Future work should focus on designing truthful mechanisms that give good approximation guarantees to the optimal social welfare and produce high social welfare in practice. The LOCALRATIO algorithm gives promising results in simulations and has good approximation guarantees and is definitely an interesting option for spectrum auctions.

A framework proposed to turn channel allocation algorithms into truthful in expectation mechanisms might be applicable, but will reduce the social welfare of the final allocation (Lavi & Swamy, 2005). A method to turn LOCALRATIO into a truthful mechanism would be very desirable.

It can be shown that GREEDYWEIGHT can produce an allocation with social welfare arbitrarily far off from the optimal social welfare. An example is when a large link has slightly higher valuation than an arbitrary number of smaller links that could form a feasible set. The large link will be determined the winner of the auction by GREEDYWEIGHT although the sum of the valuations the smaller links may be arbitrarily higher. However this special case didn't seem to come up in the simulations, except for varying valuations.

On the other hand it has been shown that LOCALRATIO can guarantee an allocation with social welfare which ρ-approximates the optimal social welfare (Ye & Borodin, 2009), where ρ is the inductive independence number of the underlying conflict graph. It has also been shown that $\rho = \mathcal{O}(\log(n))$ for wireless communication in the SINR model (Hoefer et al., 2011).

Although GREEDYWEIGHT produced high social welfare in simulations with randomly generated instances of wireless links, this does not mean that it will always do so in practice. Therefore it may be safer to use LOCALRATIO and have some guaranteed approximation.

We should be careful to make general conclusions based on simulation results since the results might be heavily influenced by the way the simulations are constructed. Whether the distribution of links in the plane, power assignment and valuations of links are realistic is questionable. Attempts should be made to overcome these shortcomings, for example by gathering information about locations of wireless devices, and using those locations instead of randomly generated once. Or that information could be used to improve our modelling of distribution of wireless links in the plane.

We limit our analysis and simulations to a special case of valuation functions, which may or may not be realistic. We assume that the valuations of links are symmetric and downward sloping. Symmetric valuations may be unrealistic for several reasons. Consecutive channels offer better transmission rates then disperse channels and links may prefer some channels to others due to hardware constraints. Downward sloping valuations may be unrealistic in some situations. Some communication may require substantial bandwidth or at least affect the quality of the communication. Some links may be willing to pay increasing amounts for additional channels and may even not be interested unless they get a minimum number of channels.

The design of approximation algorithms could take at least two directions. One would be to investigate what kind of valuation functions we should expect. Are symmetric and downward sloping valuation functions reasonable or should we expect something different? Another direction would be to design approximation algorithms which can handle arbitrary valuations. They of course have the obvious advantage of not having to make any assumptions about the valuation functions. The disadvantage of algorithms that allow arbitrary valuations is that they are probably not as efficient.

An attempt was made to design mechanisms for GREEDYWEIGHT and LOCALRATIO that are individually rational, produce no positive transfers and are truthful. A mechanism for GREEDYWEIGHT is proposed in Algorithm 3 and for LOCALRATIO in Algorithm 7. These mechanisms fulfill two of the goals, that is they are individually rational and they produce no positive transfers, but they are not truthful. The lack of truthfulness poses serious problems for an auctioneer trying to maximize social welfare. If the auctioneer cannot assume that bidders are bidding truthfully he has no idea how far off the social welfare of the allocation found is from the optimal social welfare.

It would also be desirable if the truthful mechanisms for GREEDYWEIGHT and LOCALRATIO would not need to make any changes to the allocations found when calculating payments. In Figure 4.6 we can see that the number of changes made to the allocations was not large in these simulations, but in theory we might have to change the allocation up to n times. This slows down the process of allocating channels significantly.

It might prove difficult to design a truthful mechanism for GREEDYWEIGHT since the valuation of links affects the ordering in which they are considered. It may also be difficult to design a truthful mechanism for LOCALRATIO since the valuation of links affects which links are considered in the backward pass. A link which untruthfully reports its valuation as higher than its true valuation might exclude some links from consideration that would otherwise block it from the final allocation.

The results of the simulations are quite clear, GREEDYWEIGHT and LOCALRATIO are the best options. It is tempting to conclude that the auctioneer should choose GREEDYWEIGHT, but its lack of approximation guarantee is problematic. The special case mentioned is unlikely to happen in randomly generated instances, at least as they were generated in these simulations, but we do not know whether they are unlikely to occur in practice. In the second last part of the simulations LOCALRATIO performed considerably better when we varied the exponent λ used to generate valuations. In the last part of the simulations the randomness factor in valuation generation was varied. The gap between the social welfare of LOCALRATIO and GREEDYWEIGHT was larger if more randomness is allowed (lower values of δ), in favour of LOCALRATIO. The performance of GREEDY-WEIGHT seems to depend on the way valuations are generated.

In conclusion, we recommend finding the optimal solution and applying the VCG mechanism for small instances. In our simulations we could find the optimal solution in reasonable amount of time for $n \leq 40$ and $m \leq 4$. While there is no known method of turning LOCALRATIO into a truthful mechanism we recommend applying LOCALRATIO with first price payments for large instances. Future work should focus on finding a method that turns LOCALRATIO into a truthful mechanism without losing too much of the social welfare obtained. Although the LP algorithm ranked third, it might be better suited for different kinds of valuations. Incentive compatible mechanisms with approximation guarantees have even been obtained for arbitrary valuations for the LP algorithm (Hoefer & Kesselheim, 2012).

Bibliography

Abramson, N. (1970). The ALOHA system: another alternative for computer communications. In *Proceedings of the november 17-19, 1970, fall joint computer conference* (pp. 281–285). New York, NY, USA: ACM. Retrieved from http://doi.acm.org/10.1145/1478462.1478502 doi: 10.1145/1478462.1478502

Anderegg, L., & Eidenbenz, S. (2003). Ad hoc-VCG: A truthful and cost-efficient routing protocol for mobile ad hoc networks with selfish agents. In (pp. 245–259).

Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, *286*(5439), 509-512. Retrieved from http://www.sciencemag.org/content/286/5439/509.abstract doi: 10.1126/science.286.5439.509

Brar, G., Blough, D. M., & Santi, P. (2006). Computationally efficient scheduling with the physical interference model for throughput improvement in wireless mesh networks. In *Proceedings of the 12th annual international conference on mobile computing and networking* (pp. 2–13). New York, NY, USA: ACM. Retrieved from http://doi.acm.org/10.1145/1161089.1161092 doi: 10.1145/1161089.1161092

Clarke, E. (1971). Multipart pricing of public goods. *Public Choice*, *11*(1), 17-33.

Goussevskaia, O., Oswald, Y. A., & Wattenhofer, R. (2007). Complexity in geometric SINR. In *Proceedings of the 8th ACM international symposium on mobile ad hoc networking and computing* (pp. 100–109). New York, NY, USA: ACM. Retrieved from http://doi.acm.org/10.1145/1288107.1288122 doi: 10.1145/1288107.1288122

Goussievskaia, O., Halldorsson, M. M., Wattenhofer, R., & Welzl, E. (2009). Capacity of arbitrary wireless networks. In *Proceedings 28th annual of IEEE conference on computer communications (INFOCOM)* (p. 1872-1880).

Gronkvist, J., & Hansson, A. (2001). Comparison between graph-based and interference-based STDMA scheduling. *MOBIHOC*.

Groves, T. (1973). Incentives in teams. *Econometrica*, *41*(4), 617-631.

Gupta, P., & Kumar, P. R. (2000). The capacity of wireless networks. *IEEE Transactions on Information Theory*, *46*(2), 388-404.

Halldórsson, M. M., Holzer, S., Mitra, P., & Wattenhofer, R. (2012). The power of non-uniform wireless power. *CoRR*, *abs/1210.3371*.

Halldorsson, M. M., & Mitra, P. (2011). Algorithms for weighted capacity and admission control in wireless networks. *CoRR*, -1–1.

Halldórsson, M., & Radhakrishnan, J. (1997). Greed is good: Approximating independent sets in sparse and bounded-degree graphs. *Algorithmica*, *18*, 145-163. Retrieved from `http://dx.doi.org/10.1007/BF02523693` doi: 10.1007/BF02523693

Hoefer, M., & Kesselheim, T. (2012). Secondary spectrum auctions for symmetric and submodular bidders. In *Proceedings of the 13th ACM conference on electrical commerce* (p. 657-671).

Hoefer, M., Kesselheim, T., & Vöcking, B. (2011). Approximation algorithms for secondary spectrum auctions. In *Proceedings 23rd symposium parallelism in algorithms and architectures (SPAA)* (p. 177-186).

Iyer, A., Rosenberg, C., & Karnik, A. (2006). What is the right model for wireless channel interference? In *Proceedings of the 3rd international conference on quality of service in heterogeneous wired/wireless networks.* New York, NY, USA: ACM. Retrieved from `http://doi.acm.org/10.1145/1185373.1185376` doi: 10.1145/1185373.1185376

Jain, K., Padhye, J., Padmanabhan, V. N., & Qiu, L. (2003). Impact of interference on multi-hop wireless network performance. In *Proceedings of the 9th annual international conference on mobile computing and networking* (pp. 66–80). New York, NY, USA: ACM. Retrieved from `http://doi.acm.org/10.1145/938985.938993` doi: 10.1145/938985.938993

Jia, J., Zhang, Q., Zhang, Q., & Liu, M. (2009). Revenue generation for truthful spectrum auction in dynamic spectrum access. In *Proceedings of the tenth ACM international symposium on mobile ad hoc networking and computing* (pp. 3–12). New York, NY, USA: ACM. Retrieved from `http://doi.acm.org/10.1145/1530748.1530751` doi: 10.1145/1530748.1530751

Joo, C., Lin, X., & Shroff, N. (2008). Understanding the capacity region of the greedy maximal scheduling algorithm in multi-hop wireless networks. *INFOCOM*.

Krishna, V. (2002). *Auction theory*. San Diego, USA: Academic Press.

Kumar, V., Marathe, M., Parthasarathy, S., & Srinivasan, A. (2005a). Algorithmic aspects of capacity in wireless networks. In *Proceedings of the international conference on measurements and modeling of computer systems (SIGMETRICS).*

Kumar, V., Marathe, M., Parthasarathy, S., & Srinivasan, A. (2005b). Algorithmic aspects of capacity in wireless networks. *SIGMETRICS*.

Lavi, R., & Swamy, C. (2005). Truthful and near-optimal mechanism design via linear programming. In *Proceedings 46th symposium foundations of computer science (FOCS)* (p. 595-604).

Lehmann, D., O'Callaghan, L. I., & Shoham, Y. (2002). Truth revelation in approximately efficient combinatorial auctions. *JACM*, *49*(5).

Moscibroda, T., Wattenhofer, R., & Weber, Y. (2006, November). Protocol design beyond graph-based models. *Hotnets*.

Nisan, N., Roughgarden, T., Tardos, E., & Vazirani, V. V. (2007). *Algorithmic game theory*. Cambridge University Press.

Roberts, L. G. (1975, April). ALOHA packet system with and without slots and capture. *SIGCOMM Computer Communication Review*, *5*(2), 28–42. Retrieved from `http://doi.acm.org/10.1145/1024916.1024920` doi: 10.1145/1024916.1024920

Schwartz, M., & Abramson, N. (2009). The ALOHAnet - surfing for wireless data [history of communications]. *Communications Magazine, IEEE*, *47*(12), 21-25. doi: 10.1109/MCOM.2009.5350363

Ásgeirsson, E. I., Foley, J. T., Gudmundsdottir, H., Halldórsson, M. M., Magnusson, S., Mitra, P., ... Vigfusson, Y. (2012). *Putting theory to the test: Distributed connectivity in wireless networks* (Technical report). http://syslab.ru.is/lib/exe/fetch.php?media=tr.pdf: Reykjavik University.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, *16*(1), 8-37.

Wan, P.-J. (2009). Multiflows in multihop wireless networks. In *Proceedings of the 10th ACM international symposium mobile ad-hoc networking and computing (MOBIHOC)* (p. 85-94).

Wang, W., Li, X.-Y., & Wang, Y. (2004). Truthful multicast routing in selfish wireless networks. In *Proceedings of the 10th annual international conference on mobile computing and networking* (pp. 245–259). New York, NY, USA: ACM. Retrieved from `http://doi.acm.org/10.1145/1023720.1023745` doi: 10.1145/1023720.1023745

Wu, Y., Wang, B., Liu, K. J. R., & Clancy, T. (2008). A multi-winner cognitive spectrum auction framework with collusion-resistant mechanisms. In *New frontiers in* *dynamic spectrum access networks, 2008. DySPAN 2008. 3rd IEEE symposium on* (p. 1-9). doi: 10.1109/DYSPAN.2008.9

Ye, Y., & Borodin, A. (2009). Elimination graphs. In *Proceedings of the 37th international conference on algorithms, languages and programming (ICALP)* (p. 774-785).

Yücek, T., & Arslan, H. (First Quarter 2009). A survey of spectrum sensing algorithms for cognitive radio applications. *IEEE Communications Surveys and Tutorials*, *11*(1), 116-130.

Zhou, X., Gandhi, S., Suri, S., & Zheng, H. (2008). eBay in the sky: Strategy-proof wireless spectrum auctions. In *Proceedings of the 14th international conference on mobile computing and networking (MobiCom)* (p. 2-13).

Appendix A

Simulation results

A.1 Varying ρ

ρ	μ	σ	Min	Q_1	Median	Q_3	Max
1	55.8	10.4	36.1	48.2	55.9	63.8	76.5
	48.7	7.5	33.4	43.4	49.0	55.0	65.2
2	72.7	12.3	48.9	62.5	72.0	82.1	98.3
	51.5	7.3	36.6	46.2	51.5	57.3	66.4
3	82.4	13.1	54.5	72.3	82.6	91.8	111.3
	52.0	7.0	38.9	46.2	51.9	57.2	67.1
4	88.5	13.5	58.5	77.7	88.9	98.1	119.2
	52.2	6.8	39.7	46.8	51.9	56.9	66.2
5	92.7	13.7	61.2	82.0	92.7	102.3	124.0
	52.0	6.7	39.3	47.5	51.2	57.1	66.6
6	95.5	13.8	63.1	85.1	94.3	105.3	126.9
	51.7	6.6	39.6	46.6	51.1	56.2	68.5
7	97.6	13.9	64.5	88.0	96.0	107.3	128.7
	51.7	6.4	40.2	46.5	51.4	56.0	65.6
8	99.2	14.0	65.6	89.5	98.0	108.8	130.4
	51.5	6.6	39.0	46.3	51.1	56.2	67.1
9	100.4	14.1	66.4	90.8	99.3	110.2	132.3
	51.3	6.4	39.5	46.7	51.1	56.1	64.1
10	101.4	14.1	67.2	91.6	101.1	111.1	134.6
	51.1	6.4	38.1	46.9	50.9	55.3	65.7

Table A.1: Social welfare of the solution given by the LP algorithm.

A.2 Varying Number of Links

n	μ	σ	Min	Q_1	Median	Q_3	Max
10	13.0	5.0	1.4	9.3	12.8	17.0	26.7
20	15.6	5.2	4.4	12.3	14.7	18.5	29.5
30	17.9	5.9	4.9	14.1	18.1	21.6	31.6
40	18.6	5.1	7.6	14.7	18.8	22.0	29.1
50	20.2	6.0	7.9	15.7	19.7	24.7	37.4
60	20.9	5.3	7.1	17.0	21.2	24.9	33.4
70	20.8	6.4	2.7	16.9	20.6	25.0	35.5
80	21.6	6.6	4.3	17.4	21.7	25.9	42.0
90	21.6	6.7	7.8	16.9	20.7	26.4	36.0
100	23.2	6.2	8.7	18.6	22.8	28.2	38.5

Table A.2: Social welfare of GREEDYLENGTH when varying number of links n.

n	μ	σ	Min	Q_1	Median	Q_3	Max
10	22.0	6.2	12.9	17.6	21.1	26.5	42.3
20	35.2	6.7	21.8	29.5	34.9	38.8	50.6
30	46.1	5.9	33.7	42.2	45.3	49.4	58.6
40	53.1	7.3	36.9	49.1	52.5	58.4	75.7
50	60.0	7.9	44.6	54.0	61.2	65.4	77.1
60	62.1	7.6	44.6	56.6	61.8	67.2	78.6
70	65.6	7.6	43.4	61.5	66.5	70.9	78.4
80	69.6	7.2	55.6	64.3	70.4	74.6	85.5
90	70.0	7.6	49.8	64.9	69.7	75.0	91.5
100	72.8	7.1	50.0	67.9	72.9	77.8	86.9

Table A.3: Social welfare of GREEDYWEIGHT when varying number of links n.

n	μ	σ	Min	Q_1	Median	Q_3	Max
10	16.5	5.1	8.5	13.3	15.3	19.7	35.5
20	25.6	5.9	12.6	21.3	25.1	29.4	43.1
30	32.7	5.6	21.3	29.0	32.9	36.6	45.3
40	38.8	6.5	24.0	33.7	39.4	43.2	57.2
50	44.7	6.2	28.8	40.1	44.3	49.4	59.8
60	47.1	6.9	19.7	41.7	47.3	51.4	61.9
70	48.9	7.0	33.9	44.1	49.1	53.3	68.0
80	51.3	6.1	35.4	47.3	51.3	56.0	61.7
90	52.8	6.9	36.6	48.8	51.7	57.0	66.8
100	55.2	5.5	43.0	51.5	54.8	58.7	67.1

Table A.4: Social welfare of BUCKETLENGTH when varying number of links n.

n	μ	σ	Min	Q_1	Median	Q_3	Max
10	16.0	5.1	7.9	12.1	14.8	19.3	32.5
20	23.6	5.4	13.1	20.0	22.3	27.2	35.9
30	29.5	5.4	17.6	25.7	28.8	33.1	52.3
40	35.3	6.6	18.0	30.6	35.5	39.3	58.0
50	39.5	7.1	22.8	34.8	39.6	45.0	55.0
60	42.4	6.5	29.7	37.3	41.6	46.7	61.0
70	45.1	7.1	30.3	40.6	44.8	49.3	64.2
80	47.7	6.0	33.6	43.9	47.7	50.9	63.5
90	48.8	6.2	31.9	44.9	49.2	52.5	65.4
100	49.7	6.3	38.8	44.8	49.3	54.1	65.1

Table A.5: Social welfare of BUCKETWEIGHT when varying number of links n.

n	μ	σ	Min	Q_1	Median	Q_3	Max
10	21.6	6.2	11.6	17.0	20.5	26.2	42.1
20	34.2	6.6	20.8	28.9	33.9	37.2	48.1
30	44.7	5.9	33.0	40.5	43.8	48.0	61.0
40	51.2	7.2	34.6	46.1	50.9	55.7	74.4
50	59.1	7.9	39.6	52.9	59.7	64.3	79.2
60	62.0	7.4	46.8	57.2	61.2	65.5	78.7
70	65.2	7.8	44.7	59.9	65.8	70.6	87.5
80	70.0	5.8	59.0	65.8	69.8	73.4	88.6
90	70.2	7.0	56.7	64.3	70.2	75.6	82.7
100	72.4	7.6	53.6	67.5	72.3	77.9	87.9

Table A.6: Social welfare of LOCALRATIO when varying number of links n.

n	μ	σ	Min	Q_1	Median	Q_3	Max
10	20.2	6.1	8.3	15.6	19.8	24.3	38.5
20	30.2	6.2	18.6	25.6	29.6	34.7	45.1
30	37.7	6.2	21.2	33.2	38.2	42.7	50.8
40	41.0	7.6	23.7	36.3	41.5	45.1	62.6
50	46.9	8.0	27.4	42.0	46.8	51.7	68.0
60	48.2	7.3	30.5	42.0	48.2	52.6	64.4
70	49.2	6.8	25.1	45.9	48.3	52.9	65.0
80	51.3	8.0	31.9	45.8	51.9	56.8	68.6
90	51.1	7.6	31.9	45.7	51.1	56.2	71.6
100	53.7	6.8	35.5	48.9	53.4	59.3	65.6

Table A.7: Social welfare of GREEDYINTERFERENCE when varying number of links n.

n	μ	σ	Min	Q_1	Median	Q_3	Max
10	20.7	6.8	7.3	15.9	19.7	25.6	42.1
20	32.4	6.6	20.5	28.1	32.0	35.9	49.4
30	42.1	6.4	29.5	37.6	42.1	45.7	57.5
40	47.4	7.0	28.8	43.5	48.0	51.3	67.3
50	54.5	7.5	38.8	49.0	54.9	59.2	72.5
60	58.5	6.5	45.7	53.5	57.8	63.0	73.7
70	58.9	6.1	41.0	55.3	58.9	63.3	71.1
80	63.2	6.0	51.3	58.7	62.3	67.0	76.7
90	64.0	5.9	49.8	59.0	63.9	68.6	77.1
100	66.1	5.6	54.1	62.2	65.9	70.8	76.9

Table A.8: Social welfare of LP when varying number of links n.

n	μ	σ	Min	Q_1	Median	Q_3	Max
10	22.1	6.4	12.9	17.4	21.2	26.7	42.1
20	36.2	6.9	22.7	30.9	36.4	40.5	52.0
30	48.9	6.4	34.2	43.9	48.1	54.0	64.9
40	57.8	7.7	40.0	52.5	58.2	62.6	82.3

Table A.9: Optimal social welfare when varying number of links n.

A.3 Varying Number of Channels

m	μ	σ	Min	Q_1	Median	Q_3	Max
1	9.2	3.2	2.1	7.0	9.0	11.3	17.8
2	12.5	3.4	3.4	10.9	12.9	14.8	21.4
3	16.0	4.2	5.7	13.3	16.6	19.0	25.1
4	19.8	5.1	8.5	16.3	19.6	22.8	34.3
5	21.7	7.1	2.9	16.7	21.6	26.7	37.0
6	24.4	7.3	10.5	18.6	24.7	29.3	40.4
7	28.4	9.4	9.4	21.2	28.2	36.0	47.6
8	33.2	9.5	15.3	26.4	32.9	40.4	52.5
9	37.2	10.0	12.6	29.8	36.6	44.7	65.4
10	40.9	11.8	17.0	33.5	39.9	48.9	77.3

Table A.10: Social welfare of GREEDYLENGTH when varying the the number of channels m.

m	μ	σ	Min	Q_1	Median	Q_3	Max
1	18.2	3.5	8.9	16.1	18.5	20.3	26.0
2	30.9	4.0	22.9	28.0	30.7	33.3	41.5
3	42.6	5.5	26.4	39.2	42.7	46.6	52.7
4	54.5	5.8	43.4	50.6	54.4	58.8	69.7
5	64.0	8.6	43.9	57.2	63.3	70.8	80.9
6	76.0	8.8	58.1	69.6	76.2	81.3	100.8
7	85.4	11.6	60.9	77.2	84.8	93.1	120.0
8	96.5	12.2	70.1	87.5	96.4	104.7	127.7
9	104.8	14.1	70.7	94.2	102.6	116.5	143.6
10	120.9	13.6	85.1	110.7	121.5	128.2	155.4

Table A.11: Social welfare of GREEDYWEIGHT when varying the the number of channels m.

m	μ	σ	Min	Q_1	Median	Q_3	Max
1	12.6	2.7	7.2	10.6	11.9	14.2	21.5
2	21.7	4.0	10.1	19.2	21.5	24.5	31.5
3	30.6	5.1	18.1	27.9	30.2	33.2	43.0
4	39.4	6.0	25.1	34.9	39.3	44.1	51.1
5	46.3	7.4	32.3	40.9	46.3	51.4	63.1
6	54.9	8.6	37.3	48.7	55.1	61.2	79.1
7	62.5	11.3	43.5	55.1	61.5	68.4	104.4
8	70.7	11.3	49.2	61.4	71.2	77.9	110.5
9	75.5	11.2	47.3	67.7	75.4	84.2	97.3
10	88.1	13.9	50.8	78.5	87.8	96.7	122.6

Table A.12: Social welfare of BUCKETLENGTH when varying the the number of channels m.

m	μ	σ	Min	Q_1	Median	Q_3	Max
1	11.2	3.2	5.1	9.0	10.9	13.3	21.1
2	18.3	3.5	11.1	16.0	17.9	21.1	27.6
3	27.9	5.1	14.5	24.0	28.6	30.7	37.5
4	35.7	5.8	22.3	31.6	36.1	40.3	47.6
5	41.0	7.4	25.8	34.9	42.2	46.2	58.2
6	51.2	8.3	34.4	45.5	50.4	56.8	72.6
7	56.7	9.5	36.7	49.8	56.8	62.6	79.8
8	63.7	9.6	39.5	56.9	63.5	68.9	88.4
9	68.8	11.6	44.2	60.5	67.3	77.7	99.1
10	80.4	12.7	49.7	71.9	79.9	88.0	116.7

Table A.13: Social welfare of BUCKETWEIGHT when varying the the number of channels m.

m	μ	σ	Min	Q_1	Median	Q_3	Max
1	17.7	3.6	9.2	15.4	18.1	19.8	25.6
2	30.6	4.2	21.7	28.3	30.0	33.5	41.0
3	41.7	4.9	26.8	38.6	41.2	45.8	51.4
4	54.1	6.6	39.7	50.2	54.1	59.1	69.5
5	62.9	9.0	39.8	56.8	62.5	68.8	90.5
6	74.6	9.3	50.8	68.1	74.2	81.6	95.5
7	84.1	11.1	60.3	76.7	84.7	90.1	114.9
8	94.8	12.0	67.5	85.1	94.2	104.2	126.5
9	102.2	14.4	69.7	92.9	100.1	111.7	143.4
10	120.1	15.0	85.9	107.9	120.7	128.1	160.5

Table A.14: Social welfare of LOCALRATIO when varying the the number of channels m.

m	μ	σ	Min	Q_1	Median	Q_3	Max
1	14.4	3.1	6.8	12.2	14.3	16.7	21.0
2	25.1	4.5	15.7	22.1	24.7	28.2	35.1
3	34.0	5.2	21.8	30.5	34.0	37.4	46.1
4	43.4	6.8	26.3	38.9	42.7	48.7	59.2
5	49.9	8.3	28.3	45.0	48.3	56.5	67.9
6	59.8	9.5	37.5	54.1	60.1	66.1	79.2
7	69.4	10.9	40.8	62.6	69.7	76.5	95.7
8	77.3	12.2	46.8	70.7	77.5	85.6	103.7
9	84.0	13.4	54.0	75.1	84.1	94.1	113.7
10	94.2	14.0	62.6	82.1	95.3	104.3	130.2

Table A.15: Social welfare of GREEDYINTERFERENCE when varying the the number of channels m.

m	μ	σ	Min	Q_1	Median	Q_3	Max
1	17.7	2.8	11.6	16.0	17.6	19.4	25.6
2	29.6	4.0	21.9	26.7	29.9	32.3	38.4
3	39.1	5.1	24.3	36.5	39.1	43.4	47.9
4	49.5	6.5	32.1	45.4	50.2	54.1	66.8
5	56.3	9.2	35.3	50.0	55.9	63.4	84.6
6	67.2	10.0	43.4	60.2	66.7	74.3	88.9
7	75.1	11.7	44.4	67.4	75.7	83.5	105.4
8	84.6	12.5	57.4	77.0	85.0	92.7	122.6
9	91.5	15.0	56.3	81.3	91.0	102.7	125.7
10	106.0	14.5	80.5	96.2	105.5	115.3	143.4

Table A.16: Social welfare of LP when varying the the number of channels m.

m	μ	σ	Min	Q_1	Median	Q_3	Max
1	19.7	2.2	14.1	18.2	19.4	21.1	25.5
2	34.5	4.1	26.5	31.3	35.1	36.8	45.1
3	46.5	5.9	33.1	42.0	46.3	51.0	60.5
4	57.8	7.7	40.0	52.5	58.2	62.6	82.3

Table A.17: Optimal social welfare when varying the the number of channels m.

A.4 Varying the Probability of Attaching

p	μ	σ	Min	Q_1	Median	Q_3	Max
0.0	21.3	6.3	5.2	17.7	21.9	25.8	32.7
0.2	21.6	5.3	12.0	18.0	21.7	25.4	36.3
0.4	18.5	5.2	4.6	15.5	18.9	22.3	28.1
0.6	16.4	5.8	3.0	12.1	16.3	20.5	28.2
0.8	12.8	6.0	2.7	8.7	12.1	15.6	30.0
1.0	3.2	2.5	0.4	1.2	2.4	5.1	10.2

Table A.18: Social welfare of GREEDYLENGTH when varying the probability of attaching p.

p	μ	σ	Min	Q_1	Median	Q_3	Max
0.0	56.6	6.3	43.8	51.7	56.8	60.7	72.5
0.2	56.4	7.6	43.1	50.1	55.7	62.6	76.8
0.4	53.8	7.0	36.1	49.8	54.1	58.8	70.7
0.6	48.0	7.3	30.8	42.9	47.8	53.1	65.9
0.8	40.0	9.0	20.1	32.9	40.0	44.4	65.5
1.0	19.3	2.8	12.8	17.4	19.3	21.4	28.9

Table A.19: Social welfare of GREEDYWEIGHT when varying the probability of attaching p.

p	μ	σ	Min	Q_1	Median	Q_3	Max
0.0	40.4	6.2	27.7	36.0	41.0	44.5	62.3
0.2	40.5	7.0	25.6	35.3	40.3	45.4	57.1
0.4	39.1	6.7	25.3	34.5	38.7	43.7	60.5
0.6	35.1	5.5	23.7	30.6	35.0	39.3	51.1
0.8	31.0	7.9	15.0	25.7	30.7	35.9	51.3
1.0	15.4	2.0	10.1	14.0	15.4	17.0	23.4

Table A.20: Social welfare of BUCKETLENGTH when varying the probability of attaching p.

p	μ	σ	Min	Q_1	Median	Q_3	Max
0.0	37.9	6.1	24.3	34.7	37.6	42.5	52.9
0.2	36.5	7.0	21.4	30.9	36.2	41.6	57.2
0.4	34.7	6.3	20.9	30.7	35.1	39.7	47.6
0.6	31.0	6.1	17.9	26.4	30.6	35.1	56.5
0.8	26.1	7.5	11.0	20.7	25.1	29.4	49.9
1.0	13.4	2.0	7.8	12.2	13.2	14.4	18.7

Table A.21: Social welfare of BUCKETWEIGHT when varying the probability of attaching p.

p	μ	σ	Min	Q_1	Median	Q_3	Max
0.0	55.6	6.4	37.7	51.0	55.8	60.1	73.5
0.2	55.0	7.5	41.7	49.9	54.4	60.7	75.8
0.4	52.5	7.1	34.6	48.6	51.8	58.4	67.2
0.6	47.4	7.1	30.8	42.7	47.6	51.6	64.6
0.8	39.0	8.8	18.6	32.6	38.7	45.5	57.5
1.0	18.6	2.8	12.1	16.9	18.5	20.4	28.5

Table A.22: Social welfare of LOCALRATIO when varying the probability of attaching p.

p	μ	σ	Min	Q_1	Median	Q_3	Max
0.0	46.8	6.9	30.5	42.6	46.5	50.4	62.3
0.2	45.3	7.4	27.9	39.7	44.8	51.1	64.0
0.4	41.7	7.0	27.0	36.3	41.5	46.3	60.8
0.6	36.7	6.5	20.1	32.9	36.4	41.8	50.7
0.8	29.2	7.9	10.1	23.7	28.5	33.5	53.0
1.0	10.5	3.4	2.5	8.0	10.4	13.1	20.0

Table A.23: Social welfare of GREEDYINTERFERENCE when varying the probability of attaching p.

p	μ	σ	Min	Q_1	Median	Q_3	Max
0.0	52.5	6.0	40.4	47.7	53.0	57.5	65.2
0.2	52.2	7.1	37.0	46.8	51.9	57.4	69.1
0.4	49.1	6.9	31.2	44.3	48.7	54.2	67.7
0.6	42.5	7.4	25.4	36.8	42.5	48.0	60.6
0.8	35.1	8.8	13.7	28.6	35.3	41.1	59.4
1.0	14.6	2.6	9.7	13.2	14.5	16.4	23.9

Table A.24: Social welfare of LP when varying the probability of attaching p.

p	μ	σ	Min	Q_1	Median	Q_3	Max
0.0	63.0	6.3	47.9	58.5	63.1	66.9	80.2
0.2	62.3	7.7	47.7	57.1	62.3	68.3	79.8
0.4	59.0	7.5	44.9	53.2	58.7	65.2	73.5
0.6	51.2	7.6	33.5	45.7	51.9	56.2	71.0
0.8	41.8	9.5	20.1	34.4	41.1	48.1	71.2
1.0	20.1	2.9	13.6	18.2	20.0	21.7	30.2

Table A.25: Optimal social welfare when varying the probability of attaching p.

A.5 Social Welfare and Payments for GREEDYWEIGHT and LOCALRATIO

n	Social welfare/ Total Payment	μ	σ	Min	Q_1	Median	Q_3	Max
10	Social welfare	21.3	6.1	5.7	16.9	21.4	24.3	38.5
	Total payment	4.7	2.5	0.9	3.1	4.3	5.6	11.8
20	Social welfare	36.0	7.3	17.8	32.0	35.4	40.0	54.5
	Total payment	10.3	3.7	2.5	7.5	9.8	12.7	20.8
30	Social welfare	47.5	6.6	31.3	42.7	47.5	52.4	60.6
	Total payment	15.2	5.5	0.0	11.4	15.6	19.8	29.3
40	Social welfare	56.0	6.8	39.6	51.5	55.6	60.5	74.6
	Total payment	15.5	6.9	0.0	9.5	15.9	20.0	34.6
50	Social welfare	60.2	7.0	45.5	55.2	59.1	65.6	81.5
	Total payment	13.3	7.3	0.5	7.5	13.7	18.4	33.3
60	Social welfare	66.2	6.1	49.5	63.1	66.6	70.8	83.2
	Total payment	13.5	8.1	0.0	7.1	12.6	19.3	33.0
70	Social welfare	70.3	6.1	59.4	65.4	70.6	74.1	84.7
	Total payment	10.2	7.0	0.0	4.9	9.1	13.0	33.6
80	Social welfare	72.6	7.3	56.8	67.7	71.8	77.8	87.8
	Total payment	9.6	7.3	0.0	3.7	8.9	14.3	29.9
90	Social welfare	75.6	7.1	59.6	70.2	75.1	80.6	93.0
	Total payment	10.2	7.2	0.0	4.7	8.2	15.0	28.5
100	Social welfare	77.6	5.9	62.2	73.8	78.0	80.2	93.7
	Total payment	11.4	8.5	0.0	4.3	10.1	17.5	34.1

Table A.26: Social welfare and total payment of GREEDYWEIGHT when varying number of links n.

n	μ	σ	Min	Q_1	Median	Q_3	Max
10	0.1	0.3	0.0	0.0	0.0	0.0	1.0
20	0.5	0.7	0.0	0.0	0.0	1.0	3.0
30	1.0	0.9	0.0	0.0	1.0	2.0	3.0
40	1.7	1.4	0.0	1.0	1.0	2.0	8.0
50	1.8	1.2	0.0	1.0	2.0	2.5	6.0
60	2.0	1.1	0.0	1.0	2.0	3.0	5.0
70	2.0	1.0	0.0	1.0	2.0	2.5	5.0
80	2.3	1.1	0.0	1.5	2.0	3.0	6.0
90	2.4	1.2	1.0	1.0	2.0	3.0	6.0
100	2.5	1.2	1.0	2.0	2.0	3.0	6.0

Table A.27: Changes made to the allocation of GREEDYWEIGHT while calculating payments.

n	Social welfare/ Total Payment	μ	σ	Min	Q_1	Median	Q_3	Max
10	Social welfare	20.9	6.0	5.6	16.6	20.9	24.0	37.7
	Total payment	4.7	2.7	0.0	2.9	4.3	5.9	16.4
20	Social welfare	35.3	7.3	16.8	30.8	34.8	39.6	52.2
	Total payment	8.9	4.3	0.7	5.9	8.4	10.9	22.8
30	Social welfare	46.9	6.4	34.0	42.4	46.8	51.0	58.9
	Total payment	11.3	5.7	0.5	6.7	11.1	14.6	26.0
40	Social welfare	55.3	7.0	38.1	49.9	55.1	59.8	76.7
	Total payment	12.0	6.6	0.0	7.1	11.8	16.8	31.9
50	Social welfare	59.7	7.0	44.5	55.2	59.1	64.2	76.1
	Total payment	10.9	6.0	0.0	6.0	10.1	14.6	29.4
60	Social welfare	66.0	6.1	49.8	62.0	65.7	70.7	84.4
	Total payment	10.7	6.4	0.8	6.0	9.3	14.8	29.7
70	Social welfare	71.2	6.2	55.5	66.9	71.1	75.2	88.6
	Total payment	9.4	6.9	0.0	3.9	8.3	13.3	34.4
80	Social welfare	73.0	6.9	56.2	69.7	72.5	77.5	87.4
	Total payment	9.1	5.7	0.0	4.7	8.3	13.1	22.1
90	Social welfare	77.0	7.2	62.5	71.6	76.6	83.0	94.8
	Total payment	8.3	7.1	0.0	1.8	6.2	14.1	26.0
100	Social welfare	79.4	5.8	64.6	76.2	79.2	83.1	93.6
	Total payment	8.8	6.5	0.0	3.2	8.0	12.8	25.5

Table A.28: Social welfare and total payment of LOCALRATIO when varying number of links n.

n	μ	σ	Min	Q_1	Median	Q_3	Max
10	0.3	0.5	0.0	0.0	0.0	1.0	2.0
20	0.9	0.8	0.0	0.0	1.0	1.0	3.0
30	1.6	1.3	0.0	1.0	1.5	2.0	5.0
40	2.0	1.2	0.0	1.0	2.0	2.5	7.0
50	2.5	1.1	1.0	2.0	2.0	3.0	6.0
60	2.6	1.3	0.0	2.0	2.0	3.0	7.0
70	2.7	1.4	0.0	2.0	2.5	3.0	7.0
80	2.7	1.2	1.0	2.0	3.0	3.0	7.0
90	3.1	1.3	1.0	2.0	3.0	4.0	6.0
100	3.1	1.4	1.0	2.0	3.0	4.0	7.0

Table A.29: Changes made to the allocation of LOCALRATIO while calculating payments.

A.6 Compare Valuations

λ	μ	σ	Min	Q_1	Median	Q_3	Max
0.0	13.7	4.7	5.0	10.0	13.0	16.0	31.0
0.1	8.3	1.9	4.2	6.8	8.2	9.6	13.1
0.2	9.1	2.0	4.6	7.6	9.2	10.3	13.5
0.3	10.1	2.2	5.1	8.5	10.2	11.5	14.4
0.4	11.4	2.4	5.6	9.6	11.7	13.0	16.3
0.5	12.9	2.7	6.3	11.0	13.2	14.7	18.5
0.6	14.8	3.1	7.2	12.6	15.2	16.8	21.3
0.7	17.0	3.5	8.2	14.4	17.4	19.4	24.5
0.8	19.6	4.1	9.4	16.6	20.1	22.4	28.3
0.9	22.7	4.7	10.8	19.3	23.2	26.0	32.8
1.0	26.3	5.4	12.4	22.4	26.9	30.2	38.1
1.1	30.5	6.2	14.3	26.1	31.2	35.1	44.4
1.2	35.5	7.2	16.6	30.4	36.3	40.7	51.6
1.3	41.4	8.3	19.3	35.4	42.4	47.3	60.2
1.4	48.2	9.7	22.4	41.3	49.6	55.2	70.2
1.5	56.3	11.2	26.0	48.2	57.8	64.5	81.9
1.6	65.7	13.0	30.3	56.3	67.5	75.3	95.6
1.7	76.8	15.1	35.3	65.8	78.7	87.9	111.7
1.8	89.7	17.6	41.2	77.0	91.6	102.8	130.4
1.9	104.9	20.4	48.1	90.1	107.0	119.9	152.4
2.0	122.6	23.7	56.2	105.5	125.2	140.0	178.1
2.1	143.5	27.6	65.7	123.7	146.5	163.5	208.2

Table A.30: Social welfare of GREEDYWEIGHT when varying the exponent λ used for generating valuations.

λ	μ	σ	Min	Q_1	Median	Q_3	Max
0.0	14.2	4.8	6.0	11.0	13.0	17.0	29.0
0.1	14.0	4.1	6.4	10.7	13.7	16.7	25.5
0.2	14.0	3.4	6.9	11.5	14.1	16.7	24.0
0.3	15.2	3.1	7.3	13.1	15.4	17.4	22.6
0.4	16.7	3.4	7.0	14.8	16.7	19.4	23.9
0.5	17.9	3.5	8.1	16.1	18.5	20.4	25.2
0.6	20.6	3.2	9.4	18.6	21.4	22.7	27.4
0.7	23.0	3.5	10.8	20.9	23.1	25.6	29.3
0.8	26.0	3.5	16.3	23.6	26.2	28.9	32.3
0.9	29.4	4.1	15.9	26.8	30.1	32.5	36.9
1.0	33.5	4.5	20.1	30.7	34.3	37.0	42.1
1.1	38.4	4.6	27.2	35.4	39.1	42.2	48.3
1.2	44.2	5.2	31.2	40.5	45.2	48.4	55.4
1.3	50.5	6.8	27.7	46.1	51.0	56.0	63.7
1.4	57.9	8.8	22.1	53.5	59.0	63.7	73.3
1.5	66.9	10.3	25.6	61.4	68.7	73.7	84.5
1.6	77.0	11.9	29.8	70.8	79.9	83.8	97.5
1.7	89.5	14.1	34.7	83.2	93.2	97.9	114.8
1.8	103.4	16.8	40.4	98.0	107.6	114.1	133.5
1.9	119.7	19.5	47.1	113.5	123.0	132.1	155.3
2.0	139.5	22.4	55.0	132.7	144.0	153.1	180.8
2.1	163.1	25.8	64.2	155.0	168.8	179.0	210.4

Table A.31: Social welfare of LOCALRATIO when varying the exponent λ used for generating valuations.

δ	μ	σ	Min	Q_1	Median	Q_3	Max
0.0	108.4	11.4	77.0	100.5	108.2	116.0	140.8
0.1	105.0	10.2	68.5	97.8	106.1	112.1	123.4
0.2	100.1	9.7	81.1	92.9	98.3	107.4	129.8
0.3	98.2	10.4	69.2	91.4	99.6	104.5	117.5
0.4	94.9	9.8	73.5	88.0	95.2	100.8	124.1
0.5	93.3	9.3	71.6	86.5	93.2	99.1	125.4
0.6	91.3	8.7	70.1	85.8	90.8	97.4	111.6
0.7	90.3	10.5	66.6	83.0	91.0	98.6	118.9
0.8	88.2	8.1	70.7	83.0	87.1	93.6	106.7
0.9	85.8	9.0	61.3	80.3	86.8	92.4	104.9
1.0	87.8	7.6	70.0	83.1	87.3	93.1	104.9

Table A.32: Social welfare of GREEDYWEIGHT when varying the randomness factor δ used for generating valuations.

δ	μ	σ	Min	Q_1	Median	Q_3	Max
0.0	123.0	10.5	99.3	115.6	124.7	130.7	141.5
0.1	116.6	9.8	95.6	107.9	116.6	124.4	138.3
0.2	112.7	8.9	90.8	106.5	113.1	119.5	133.6
0.3	108.6	8.3	91.9	102.9	107.9	114.3	136.7
0.4	103.3	8.1	84.9	98.7	103.5	107.2	120.3
0.5	99.4	9.1	64.3	94.1	100.1	104.7	117.5
0.6	97.6	7.2	79.7	93.0	98.1	102.5	114.4
0.7	95.1	8.5	73.4	89.5	94.4	100.6	119.6
0.8	90.4	8.2	69.8	84.5	90.3	96.1	108.6
0.9	89.0	10.6	45.2	82.9	88.4	95.6	112.2
1.0	89.9	7.9	72.7	83.3	91.0	96.0	113.7

Table A.33: Social welfare of LOCALRATIO when varying the randomness factor δ used for generating valuations.

School of Computer Science
Reykjavík University
Menntavegi 1
101 Reykjavík, Iceland
Tel. +354 599 6200
Fax +354 599 6201
www.reykjavikuniversity.is
ISSN 1670-8539