On the Benefit of Cooperation of Secondary Users in Dynamic Spectrum Access

Justin M. Kelly

Thesis submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Master of Science
in
Electrical Engineering

R. Michael Buehrer, Chair
Claudio da Silva
Allen B. MacKenzie

July 27, 2009
Blacksburg, Virginia

Keywords: Cognitive Radio, Dynamic Spectrum Access, Power Control, Dynamic Spectrum Sharing, Rate Adaptation

Benefit of Cooperation of Secondary Users in Dynamic Spectrum Access

Justin M. Kelly

ABSTRACT

For the past 70 years, the Federal Communications Commission (FCC) has been the licensing authority for wireless spectrum. Traditionally, spectrum was commercially licensed to primary users with defined uses. With the growth of personal communication systems in the 1990's, unallocated spectrum has become a scarce commodity. However, since most primary users are active only at certain times and places, much of the allocated spectrum remains underutilized. Substantial holes exist in the spatio-temporal spectrum that could be opportunistically used by unlicensed secondary users. As a result, the FCC is considering allowing secondary users to opportunistically use frequencies that are not being used by primary users. If multiple secondary users are present in the same geographical area, the concept of Dynamic Spectrum Sharing (DSS) allows these users to share the opportunistic spectrum.

If several secondary users want to use a limited set of frequency resources, they will very likely interfere with each other. Sensing is a distributed technique where each transmitter/receiver pair senses (both passively and actively) the available channels and uses the channel that provides the best performance. While sensing alone allows sharing of the spectrum, it is not the optimal method in terms of maximizing the capacity in such a shared system. If we allow the secondary users to collaborate and share information, optimal capacity might be reached. However, collaboration adds another level of complexity to the transceivers of the secondary users, since they must now be able to communicate (Note that in general, the secondary users may have completely different communication protocols, e.g., Wi-Fi and Bluetooth). Additionally, optimizing the capacity of the available spectrum could have other negative side effects such as impacting the fairness of sharing the resources. Our primary goal is to explore the benefit of this cost-benefit tradeoff by determining the capacity increase obtainable from collaboration. As a secondary goal, we also wish to determine how this increase in capacity affects fairness. To summarize, the goal of this work is to answer the question: Fundamentally, what is the benefit of collaboration in Dynamic Spectrum Sharing?

Acknowledgments

I would like to mostly thank Dr. Buehrer for all the time he invested into my education and progress over the past two years. Without this, I would have never been able to do any of this work. Thanks also to Dr. DaSilva and Dr. MacKenzie for participating on my committee and giving me feedback on this work. I would like to thank the professors from which I have taken classes and gained a great deal of knowledge about wireless communications.

I would also like to thank my parents for encouraging me to continue my education and supporting me through the past two years. I would especially like to thank my fiance, Melissa, who continuously encouraged me every day to keep working hard. Many times when I felt discouraged and felt like quitting, she encouraged me to keep working hard.

I would also like to thank many of the people in the lab that both helped me with my research and made my time in Blacksburg enjoyable. To Jesse, I'm glad I was able to get to know you during my time at Virginia Tech. I wish you the best in your marriage and hope I will be able to stay in touch with you as we continue through life. To Chris, I enjoyed getting to know you and the lunches we were able to share. To Harris, I enjoyed the lunches we shared discussing everything from Cyprus to the stock market. To Harpreet, thanks for all the fun conversations we shared in our cubicle about various topics. I wish you the best as you finish your degree continuing this work. To Tao, thanks for helping me to get started in research during my first year when we shared an office. To Dinesh, I'm glad I've gotten to know you better over the past year. I hope we stay in touch over the years. Thanks to Chris Phelps and Joe Gaeddert for helping me in general knowledge of MPRG and other processes (such as this template).

Thanks to Cindy Hopkins, Hilda Reynolds, and Nancy Goad for help in the many administrative tasks that they helped me with during my time at Virginia Tech. There are many other people I may not have been mentioned explicitly, but I would also like to thank these people for their help.

List of Abbreviations

BPC	Binary Power Control
DSA	Dynamic Spectrum Access
FCC	Federal Communications Commission
FPC	Full Power Control
ISM	Industrial, Scientific and Medical
MC	Multi-Channel
MUI	Multi-User Interference
NPC	No Power Control
QoS	Quality of Service
SC	Single-Channel
SDR	Software Defined Radio
SINR	Signal to Interference and Noise Ratio
SNR	Signal to Noise Ratio
SNR_{min}	Minimum Signal to Noise Ratio
SS	Spread Spectrum
UWB	Ultra Wide-Band

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The first radio receiving and transmitting device was invented by Marconi in 1897. From early on, the US government recognized the value of wireless spectrum, passing laws regulating its use as early as 1912 [3]. With the passing of the Communications Act of 1934, the government created the Federal Communications Commission (FCC) and gave it full power to oversee and license commercial use of the wireless spectrum.

Originally, the FCC allocated spectrum by separating it into a set of discrete frequency bands and assigning those bands based on extensive spectrum planning. Licenses were assigned in a "'beauty contest' mode which involved significant lobbying to regulation authorities" [4]. With the rise of personal communication systems (PCS) and other wireless technologies in the past 20 years, this planning became almost impossible causing the FCC to change its spectrum allocation technique to a market-based approach. Instead of planning out future usage, the FCC elicited input from industry on how to allocate and assign frequency bands and spectrum auctions were introduced.

The drastic rise in wireless spectrum usage in the past 20 years combined with the limited spectrum has made unassigned wireless spectrum a scarce commodity [5]. However, several studies have noted that while there is a scarcity of unlicensed spectrum, much of the licensed spectrum is underutilized [6, 7]. In an attempt to allow for higher spectrum utilization, two concepts have emerged to enable unlicensed use of the licensed spectrum, namely spectrum underlay and spectrum overlay [8].

For both systems, a user with a license to use a frequency band is called a primary or legacy user, while an unlicensed user using either an underlay or overlay method to access the licensed frequency bands is called a secondary user. The underlay method essentially spreads the power of the secondary user transmission over such a wide frequency bandwidth that

the interference power seen by any primary user is negligible. The most common technology that implements an underlay technique is ultra-wideband (UWB) communications (see [9] for more information on UWB). For overlay systems, secondary users avoid causing interference to primary users by transmitting in the 'white spaces' or underutilized frequency bands of the spectrum. The overlay system concept is often termed Cognitive Radio or Dynamic Spectrum Access (DSA). Some hybrid systems have emerged that combine the concepts of underlay and overlay by sensing the spectrum but also using a wide bandwidth [10], and it has further been suggested that these hybrid systems might allow superior performance over the overlay and underlay schemes [11].

In DSA, if several secondary users want to use a limited set of frequency resources they will very likely interfere with each other. Sensing is a distributed technique where each transmitter/receiver pair senses (both passively and actively) the available channels and uses the channel that provides the best performance. While sensing alone allows sharing of the spectrum, it is not the optimal method in terms of maximizing the capacity in such a shared system. If we allow the secondary users to collaborate and share information, optimal capacity might be reached. However, collaboration adds another level of complexity to the transceivers of the secondary users, since they must now be able to communicate (Note that in general, the secondary users may have completely different communication protocols, e.g., Wi-Fi and Bluetooth). Additionally, optimizing the capacity of the available spectrum could have other negative side effects such as impacting the fairness of sharing the resources. Our primary goal is to explore the benefit of this cost-benefit tradeoff by determining the capacity increase obtainable from collaboration. As a secondary goal, we also wish to determine how this increase in capacity affects fairness. To summarize, the goal of this work is to answer the question: Fundamentally, what is the benefit of collaboration in Dynamic Spectrum Sharing?

1.1 Literature Review

In our work, we study the overlay system approach, particularly the application of Cognitive Radio and Dynamic Spectrum Access (DSA). Because new regulation will be necessary to enable DSA users to access the 'white spaces' in the spectrum, we begin by introducing the three most popular suggestions for DSA regulation. We then briefly introduce cognitive radio and explain the technical challenges involved. We will conclude by reviewing previous work on spectrum sharing in DSA, the main focus of our work.

1.1.1 Types of Regulation

Three models have been suggested to regulate how unlicensed secondary users are allowed to access frequency bands. These include the exclusive use model, the shared use or hierarchical model, and the commons model [4, 12].

Exclusive Use Model

The exclusive use model gives the licensed owners the property rights to the spectrum. In this model, spectrum owners have the option to auction off blocks of their own unused or underutilized spectrum to secondary users on either a short-term or long-term basis. Secondary users would compete in these auctions to bid for the use of frequency bands. A considerable amount of research has studied the auction system and bidding strategies (i.e., [13, 14]). However, this model would require primary users to develop and deploy auction systems. Furthermore, all users would need to share a common protocol to negotiate spectrum usage and pricing, adding to the complexity of the system. It would also require secondary users to be within range of an auction system to be able to use secondary spectrum, which would prevent full spectral efficiency.

Shared Use or Hierarchical Model

The shared use model allows users to freely and dynamically access the underutilized spectrum with only regulation governing its use of currently licensed spectrum. Essentially, the cognitive or secondary users would sense the spectrum to find underutilized spectrum and transmit in the 'white spaces'. The secondary user would be required to avoid causing harmful interference to the primary user both by reliably sensing the presence of a primary user in the spectrum initially and vacating frequency bands upon transmission of a primary user. The main challenge to the shared use model is creating methods to reliably sense the spectrum and determine of the presence of a primary user.

Commons Model

The commons model gives all users equal access to the full spectrum (thus licensing becomes irrelevant). The main technical challenge with this approach involves developing a spectrum access technique that avoids the 'tragedy of commons' [4], that is scalable for large networks, and that is able to punish 'cheaters'. Further, it is likely that current license holders would strongly oppose the loss of the licenses that grant them exclusive or prioritized use to spectrum (especially those license holders with currently deployed systems).

An alternative 'private commons' approach similar to the shared use model is mentioned in [4]. The private commons approach is similar to the shared use model in that only regulation governs the secondary use of licensed spectrum; however, unlike the shared model, the regulation for each frequency band of the private commons approach is created by the primary license owner and thus may differ between frequency bands. The private commons model would require that the secondary user keep detailed information about the regulations on each frequency band, which would add to the complexity of the secondary user. It also does not require that the primary users be willing to share with secondary users, which may prevent full spectral efficiency; however, the private commons approach would cater the secondary user access rules to the use of the channel. For example, this could minimize interference to the primary user by requiring that the secondary user transmit a signal with specific characteristics which minimize interference on that particular channel.

Comparison of Regulatory Techniques

While several papers exist supporting each the exclusive use model [15, 16, 17] and the spectrum commons model [18, 19], the challenges of the hierarchical model seem less difficult to solve. The exclusive use model would require the implementation of spectrum servers and would require users to always be in contact with these spectrum servers adding several levels of complexity. Since primary users are not required to share spectrum and secondary users must be within range of a spectrum server to use secondary spectrum, it is not guaranteed to provide much benefit to spectrum utilization. The spectrum commons has the challenge of developing a spectrum access technique to enable it to run smoothly. The spectrum commons model would also be difficult to implement because it would be strongly opposed by current spectrum license holders (who would want to maintain their licenses for exclusive or prioritized use). The main challenge of the hierarchical model is reliably avoiding causing interference to the primary user (a significant amount of research has already considered this topic). The complexity of the exclusive use model and the resistance to the spectrum commons model are likely to make the hierarchical model the regulation of choice. In this work, we assume the hierarchical model of regulation.

1.1.2 Cognitive Radio

The concept of cognitive radio emerged from a technology called software defined radio (SDR). Compared to traditional transceivers which are generally static and single purpose, SDR transceivers are able to dynamically adapt wireless communication parameters (i.e. modulation type, frequency, coding scheme, etc.) by simple changes in software (for more information on SDR see [20]). From this idea, Joseph Mitola III and Gerald Maguire Jr. suggested that we could allow these SDR transceivers to recognize their own settings and

adapt their parameters automatically to achieve desired results (or become cognitive) [21].

Shortly after the original paper on cognitive radio, Mitola wrote a second paper suggesting a new application of cognitive radios that he called spectrum pooling [22] from which dynamic spectrum access (DSA) emerged. The general concept behind DSA is that secondary users opportunistically use frequency bands that are underutilized by primary users. DSA requires many new technical challenges to be solved before it can be adopted.

Two rules govern how secondary users must act when opportunistically using licensed bands. First, the secondary user much avoid causing harmful interference to the primary user by appropriately sensing the channel and selecting a frequency band in which to transmit. Second, the user must vacate the band immediately upon the return of the primary user. These two rules along with the desire of maximizing spectrum efficiency among secondary users leads to four main areas of research for DSA: spectrum sensing, spectrum decision making, spectrum sharing, and spectrum mobility [23].

Spectrum Sensing

Spectrum sensing involves wireless receivers sensing the channel and processing the sensing data to determine which frequency bands are unused by primary users and can be used opportunistically. Spectrum sensing is the most important technical challenge to the implementation of cognitive radio because reliable spectrum sensing is necessary to minimize harmful interference to primary users [24]. The two main approaches that a wireless receiver can use for spectrum sensing are energy detection and feature detection [23]. Some research in spectrum sensing has focused on the benefit of cooperative spectrum sensing (allowing wireless receivers to share sensing data to improve the spectrum sensing reliability by avoiding the hidden node problem [25]).

Energy detection, the simplest spectrum sensing technique, involves determining the presence of a signal by comparing the energy found within some frequency band to the noise floor. Energy detectors perform poorly when the strength of the detected signal is not high compared to the power of the noise floor. While many feature-based detection techniques exist, the main techniques include matched filtering, waveform-based sensing, and cyclostationarity-based sensing. Matched filtering and waveform-based techniques require some prior knowledge about the signal to work effectively. Matched filtering is the optimal technique for detection of primary users if the features of the signal such as modulation type and transmission frequency are known [24]. Waveform based techniques work by testing the signal for an expected pattern; for example, in ATSC transmission the PN-511 sequence (a fixed pattern transmitted every 24.2 ms.) can be used to determine the presence of a digital

TV channel [23]. Cyclostationarity-based sensing involves determining the statistics of the received power to distinguish between noise and the presence of a signal. An overview of current research in spectrum sensing can be found in [23, 24].

Two main models that exist for spectrum sensing in hardware include the single-radio and dual-radio implementations. In single-radio implementation, time slots are devoted to spectrum sensing; however, with dual-radio implementation sensing can be performed in parallel with channel usage given that sensing is performed on channels not being used by the secondary user. For the more common single-radio implementation (more common because it is less expensive and more practical), devoting time to spectrum sensing decreases spectral efficiency. Further, a tradeoff exists between the reliability of the spectrum sensing and the sensing time required. Therefore a balance between sensing spectrum reliably and maximizing secondary spectrum efficiency is necessary [24]. For this study, we assume that we have perfect spectrum sensing (i.e., no primary users exist in the frequency bands we use).

Spectrum Decision Making

Spectrum decision making involves determining which of the unused frequency bands resulting from the spectrum sensing step is the best to use. This includes knowing the statistics of the primary user (i.e., access probabilities or interference seen at the primary user), the path loss, and wireless link errors for each frequency band being considered [26]. Based upon the parameters of the given set of unused frequency bands, a selection is made based upon the requirements of the channel.

Spectrum Sharing

Spectrum sharing is the coordination between secondary users (whether implicit or explicit) to share spectrum resources efficiently. Spectrum sharing can either be performed by a centralized algorithm or a distributed algorithm, and the users can either be cooperative or non-cooperative [26]. Because spectrum sharing is the main focus of this work, we do a more thorough review of spectrum sharing techniques in Section 1.1.3.

Spectrum Mobility

If a primary user returns to a channel that secondary users are occupying, the secondary users are required to vacate the channel and find other spectrum opportunities. Because secondary users are never guaranteed to have a reliable channel, spectrum mobility allows users to dynamically switch frequency bands. The frequency bands are unreliable for two reasons: primary users may return at any time and require secondary users to vacate the

band, and if the secondary user is moving the spectrum opportunities will change [26]. Another issue that can arise in spectrum mobility is the rendezvous of the secondary users if they are forced to vacate a frequency band due to the return of the primary user (i.e., [27]). Secondary users can avoid this by using multiple non-contiguous channels [26].

1.1.3 Performance of Secondary Users (Spectrum Sharing)

Allowing secondary users to obtain good spectral efficiency has been a major research topic in recent years. Because cognitive radio is a relatively new concept, the research that studies it contains a wide variety of modeling techniques.

In general, all modeling techniques use a common set of parameters. We will begin this section by introducing these parameters and the values that these parameters can have. After this, we will introduce two concepts (fairness and security) that are important when considering the value of a research work. In the final subsection, we will introduce individual research papers and describe the relevance of each to our work.

Objective Functions

Numerous objective functions exist that all intend to 'optimize performance'. Because the number of objective functions that exist is so large, we will not introduce the objective functions here; instead, we will introduce the optimization function for each individual work in the final subsection. In general, objective functions tend to either optimize some measure of capacity, some measure of fairness, or some balance between these.

Listen-before-Talk vs. Treating Interference as Noise

In a multi-user network, interference is an important consideration in the system model. Most research works either assume that a listen-before-talk (LBT) technique is used or interference is treated as Gaussian noise. The main difference is that in the LBT technique each channel is occupied by a maximum of a single user at a time, whereas when we treat interference as noise, multiple users may transmit simultaneously on the same channel. Most LBT techniques share the channel over time by assuming that users have bursty transmission. In [28], a comparison of listen-before-talk and considering interference as noise is presented. It is shown that some users benefit from listen before talk while other users benefit from considering interference as noise. Listen before talk becomes more favorable as the number of users sharing the channel increases.

Interference Model

When interference is treated as noise, two models of interference can be used, namely the protocol model or the physical model [1]. The protocol model requires that two conditions be met for a transmission to be successful. First, the intended transmitter must be within a certain distance of the intended receiver, and second, no other transmitters within a possibly different fixed distance of the intended receiver may be transmitting on the same channel. The protocol model is relatively simple to simulate, but neglects shadowing and fading. It determines whether some threshold capacity is likely to be met but does not determine the exact capacity of the link.

On the other hand, the physical model determines the exact received powers from all transmitters and the noise power at the intended receiver and determines the signal to interference and noise ratio, $SINR$, to determine the exact capacity that can be achieved. The physical model includes the shadowing and fading that the protocol model neglects, but adds computational complexity to the simulation.

For our simulations, we assume the physical model. We determined this model to be more beneficial because it more accurately models the shadowing and fading that are inherent to a real channel. Furthermore, since our goal is to optimize capacity and not the number of admitted users, the physical model is necessary to determine the achievable capacity of a given scenario.

Collaboration and Centralized Decision Making

Another modeling separation is whether different secondary users communicate, and if these users do communicate, whether they make decisions in a distributed or centralized way [26]. We will separate these concepts by calling networks collaborative or non-collaborative if they do or do not communicate with each other and distributed or central if the spectrum decisions are made in a distributed or central manner. For this work, we intend to compare non-collaborative with collaborative techniques. Our collaborative techniques will assume a central decision since this is the most optimal situation and will show us the maximal benefit of collaboration.

Regulatory Type

Previously, we mentioned several regulatory techniques that have been suggested to make DSA a plausible option. In general, most research papers in secondary spectrum sharing focus

on either the exclusive use or the hierarchical model. Exclusive use papers tend to focus on the design of auctioning systems to assign spectrum to secondary users. Hierarchical model papers tend to focus on extending results that are optimal in the collaborative centralized approach to the collaborative distributed or the non-collaborative approaches.

Presence of Primary Users

Several variations exist on the presence of primary users in current research. Some papers assume that bursty primary users may return at any time, while some papers assume that primary users are not present. Some papers assume that the available channels differ between secondary users due to primary users being in range of some secondary users but not others. This work assumes primary users are not present, and that all secondary users have the same set of available channels (such as is the case in the T.V. 'white spaces' for a set of geographically close secondary users).

System Architecture

Different network architectures have been considered for the use of DSA. The main divisions include multi-hop, single-hop, and network-centric. The main differences between single-hop and multi-hop is that in a single-hop network the data arrives at the intended receiver after only a single transmission, whereas in a multi-hop network the data may be transmitted multiple times by multiple nodes until the data arrives at the intended receiver. Optimization over multi-hop networks involves the added complexity of routing (since multiple paths can be taken from transmitter to receiver). A network-centric architecture focuses on the situation in which nodes communicate with a common base station such as is the case in a cellular network. A single-hop architecture can be even further split into the paired node architecture and the node cloud architecture. With paired nodes, each node communicates with only a single other node but with clouds of nodes each node may communicate with as many other nodes as desired. In this work, we assume single-hop with paired nodes.

Fairness

When considering previous work in DSA, we need to consider a few factors that some, but not all authors address. While many techniques try to optimize total throughput of the system, it has been shown that optimizing total throughput is at odds with providing fairness [29] (good throughput for all users). While some research considers fairness [17, 30], there is a considerable amount that does not mention fairness (i.e, [31]). It is important for us to consider how fair the 'optimal' result is in such cases.

Security

Another important consideration when reviewing spectrum sharing research works is preventing secondary user 'cheating' in attempts to obtain better capacities or longer time slots. In [32], two examples of networks exposed to malicious attacks are given. The author suggests that because of the possibility of malicious attackers, security should be a design goal for DSA. In [33], a game-theory based technique is suggested that enforces honesty and punishes 'cheating' users. For the collaborative techniques we consider in this report, we assume that all users are perfectly honest.

Individual Works

LBT-based techniques
Because LBT techniques require users to share the channel over time, we will first briefly study design of medium access control (MAC) protocols. MAC protocols define rules that govern how users may access the channel(s). MAC protocols can be divided into two groups: contention-based and non-contention-based [34]. The most common contention-based protocols are ALOHA [35], Slotted ALOHA [35], and CSMA [36]. The most common non-contention-based protocols include Reservation ALOHA [37], and MACAW [38]. Important considerations in wireless MAC design include the hidden node problem and the exposed node problem which are more thoroughly explained in [39].

In [40], a busy-burst in a minislot is suggested as a solution to the hidden and exposed node issues in DSA networks. While the busy burst is suggested to cause about 10% overhead, it is possible that communication could be incorporated into this busy burst to reduce this overhead.

In recent years with the rise of multi-channel communications, the design of multi-channel MAC protocols has been a topic of research. A good review of multi-channel MAC protocols can be found in [41]. This study divides these protocols into the following four classes: dedicated control channel, split phase, common hopping, and McMAC. With each of these, the challenge is to separate the control messages and data messages in such a way to prevent collisions and maximize spectral efficiency. The comparison and simulation results of these four techniques in [41] show that McMAC and the common control channel tend to have the best performance in most scenarios; however, the common control channel technique requires users to have two transceivers.

In [42], a dynamic distributed channel assignment scheme called Distributed Adaptive Channel Assignment (DACA) for mesh networks based upon a simple three step algorithm is proposed. However, this algorithm is focused around the multi-hop architecture and would

not have an equivalent algorithm in a single-hop architecture. This algorithm is shown to perform better than the Slotted Seeded Channel Hopping scheme shown in [43] in the presence of a dynamic channel availability. In [43], SSCH is proposed as sharing technique to improve the capacity of IEEE 802.11.

In [44], the contention-based protocols of the newly lightly licensed 3.65-3.7 GHz band are described for the case of sharing between WiMAX and Wi-Fi. These protocols assume that networks know what other networks are present since announcing presence is required for fixed access points operating in this band. This analysis shows that equal time sharing is possible when multiple networks are aware of the presence of neighboring networks and use protocols to share the channel fairly.

In [45], a routing algorithm for a multi-hop network is analyzed using a layered graph topology. The algorithm is shown to produce fewer dropped packets than a 'Sequential Assignment' algorithm. This study is based upon the assumptions of using listen-before-talk and different users having different available sets of channels.

A group of research works have focused on using a Markov chain model to model the optimal channel access probabilities to achieve fair channel access time. All of these works assume the listen-before-talk technique.

In [46], a special scenario of a Markov chain model is considered as a model for Dynamic Spectrum Sharing to determine optimum steady state probabilities for spectrum access. It is shown that a Homo-Equalis society model based access scheme achieves near optimal results for the specific scenario considered (the Homo-Equalis society model has the benefit of significantly decreasing the cooperation overhead required). This work is made slightly more generic in [47] to include scenarios where the spectrum overlaps in differing measures called complete partitioning, complete sharing, and virtual partitioning. It is further extended in [48] to consider the case where one secondary user is prioritized over the other. This analysis shows that prioritization can be performed with the combination of special frequency permissions and Markov modeled access probabilities.

Optimal throughput is studied through optimizing the channel access probabilities of a Markov chain model in [49]. Proportional fairness is also considered as a means of balancing fairness against total capacity while max-min fairness is almost neglected because of the degradation it causes to the highest capacity channels. However, this analysis assumed there is a secondary user base station the coordinates the channel access probabilities and only assumes the presence of a single channel that secondary users access. The work is extended in [50] to remove the secondary base station and allow the secondary users to optimize their channel access probabilities based on local information.

In [51], a Markov model is analyzed that models primary users as narrowband users and secondary users as wideband users. It shows that by adjusting the density of the secondary users and the transmission probability, the QoS of the primary users can be maintained while allowing secondary access to the spectrum.

Protocol Model
In [52], five rules are suggested as a distributed way to balance capacity and proportional fairness in WiFi and WiMAX nodes. This analysis assumes the protocol model of interference and uses a centralized network architecture. Simulation shows that this rule based approach performs almost as good as collaborative approaches proposed in [30]. It compares the spectrum allocation problem to a variant of the graph coloring problem and analyzes all options. This study uses proportional fairness as the optimization metric; however, threshold capacity is used to calculate the proportional fairness metric since the protocol model is used. The rules are reiterated in [53] where they are also applied to a single-hop DSA network. Because these techniques use the protocol model, the validity of the solution being proportionally fair is uncertain for a real situation.

In [54], a graph theoretic centralized algorithm and cooperative distributed algorithm are introduced. They are introduced for both a paired single-hop and a centralized architecture based upon the assumption of using the protocol model. The algorithms introduced are focused on optimizing the channel selection to minimize the number of necessary channels for all users to achieve the threshold capacity. No simulation is used to verify the performance of either. This analysis is not particularly relevant to us because it uses the protocol interference model and it minimizes number of channels instead of maximizing number of admitted users given a fixed number of channels (as would be the case in a real scenario).

In [55], a modified protocol model is suggested to model a multi-hop spectrum sharing scenario. This analysis introduces a simplification of the branch and bound algorithm generally used to solve mixed-integer non-linear optimizations. The same authors extend their work in [56] by developing a distributed algorithm that achieves near optimal capacity for routing, channel allocation and power control. This algorithm uses a unique 'scaling factor' approach to distributed optimization that seems as if it would require significant overhead. This approach is based upon the protocol model and thus is not largely relevant to our work.

In [57], a multi-hop spectrum sharing scenario is analyzed using the protocol model for interference and an island Genetic Algorithm. The island Genetic Algorithm is meant to be an extension of the previous work by the same authors in [58] where the generalized genetic algorithm was used. The island Genetic Algorithm reduces the amount of information necessary at the node and allows it to make decisions based on local information preventing the

complexity of the algorithm from scaling exponentially with the size of the network.

In [59], a multi-hop scheduling algorithm is used to optimize max-min fairness (which consequently achieves near optimal throughput). This scheduling algorithm is based upon the multi-hop architecture and assumes the protocol model, so it is mostly irrelevant to our work.

Physical Model
In [60], dynamic channel allocation by a central access point to maximize SIR for individual users is studied. This analysis is not particularly relevant to our work because its focus is the dynamic channel allocation aspect.

In [61], a DSA approach where nodes use CDMA to communicate with base stations is considered. Users are assumed to use non-cooperative, distributed approaches to selfishly select frequency bands. A game theoretic approach is proposed to model this interaction and is compared to other strategies. The special case scenario mentioned is not relevant to our work, but the game theory modeling approach might be of value.

In [62], a distributed learning algorithm for channel selection for WLAN called Communication Free Learning (CFL) is explained. CFL is intended to avoid interference caused by multiple base stations being on the same channel by converging to a nearly interference-free equilibrium. A more detailed description of CFL can be found in [63]. This simulation focuses mainly on spectrum sharing between centralized networks; however, the algorithm might also be beneficial to non-centralized networks.

In [64], a distributed algorithm using collaboration is introduced for power control, routing, and channel allocation in multi-hop cognitive mesh networks using an interference temperature constraint. Further, this paper introduces a number of metrics related to multi-hop networks. No analysis or simulation is used to determine the benefit of using this algorithm so its value is somewhat unknown.

In [65], the GADIA algorithm is studied as an adaptive distributed method for optimizing capacity in multi-hop cognitive radio networks. The GADIA algorithm groups nodes into clusters where only one link in the cluster may be active at a time and optimizes by selecting the channel that minimizes interference seen at the receiver. While this algorithm may minimize interference seen at the receiver, it is not guaranteed that this will optimize any other performance measure. Furthermore, the design of this algorithm is focused around the multi-hop architecture, which makes it irrelevant to our work.

In [66], a game theoretic scenario is considered where primary and secondary users are both

modeled as players in the game using a game model called the Game theoretic CR Spectrum Sharing Algorithm (GCSS). This analysis in not particularly useful because it models the primary users as being part of the game.

In [67], a game theory based approach of optimizing rate and channel selection to allow a certain QoS for all users while minimizing power is considered. This analysis considers a centralized graph coloring approach, a distributed cooperative approach, and a distributed non-cooperative approach. This analysis is not particularly relevant as the optimization is intended to minimize power.

A game theory approach to jointly optimize rate and power control is presented in [68] for the case of CDMA networks that all share a common channel. This analysis uses a rate based utility function. This is not particularly relevant to our work because it requires all users to be spread spectrum which may not be realistic.

In [69], a constrained optimization problem approach is used where all users use spread spectrum techniques across all channels equally (essentially making the optimization problem only over power control). The constraints include an interference temperature and minimum QoS SINR required for the secondary users. In the case where all users can have the QoS constraint met, the optimization routine can greatly be simplified by using a geometric programming method. For the case in which not all QoS constraints can be met, a joint coordination and power control algorithm requiring limited local information is introduced and is proven iterate to a Nash equilibrium. They extend this work in [70, 71] to consider the case where secondary users might be prioritized. Neither of these analyses consider security. A similar analysis is done in [72] where QoS minimum exists for a spread spectrum cellular network with primary and secondary users. In this paper, two centralized algorithms are analyzed with average and instantaneous (including fading) channel gains and fairness is also considered.

Mechanism design is considered in [73] as a method for using game theory and being resistant to cheating nodes. Deterrents are introduced by making the performance of the secondary user worse if the secondary user attempts to cheat for its own benefit. This analysis has two parts: one that assumes unlicensed secondary sharing and one that assumes a central spectrum server that auctions spectrum.

Optimal secondary user bandwidth is considered in [74] specifically considering how many channels are optimal for the secondary user that uses multiple channels (since using more channels requires more overhead due to spectrum sensing and spectrum mobility but gives greater throughput). However, the analysis for a multi-secondary user case assumes a central assignment scheme that does not allow users to overlap and does not give consideration to

the situation of having few channels and several secondary users.

In [75], a game theoretic framework is used to compare a selfish utility function with a cooperative utility function. This analysis uses a very similar system model to our own. It is shown that the cooperative utility function (which is designed to consider fairness) is capable of increasing capacity and fairness.

In a similar work in [31], an asynchronous distributed pricing scheme is tested for its ability to yield better average capacity per user for the case of a single channel limit per user. This analysis compared the limited cooperation approach where users share interference costs only and each pair bases their decisions on a utility function that combines selfish capacity and interference costs of neighboring users. With this minimal cooperation, it is shown that a benefit can be achieved over distributive approaches. It is shown that continuous power control performs better than no power control in these simulations (being able to turn users off benefited the system throughput). In [76], the asynchronous distributed pricing scheme is extended for the case of users capable of using multiple channels. For certain cases, the Asynchronous Distributed Pricing Scheme can be proven to converge and is optimal. This work however neglects to mention fairness and assumes that all users are honest when sending pricing information.

In [77], a game theoretic approach that uses punishment strategies to enforce security is suggested as a way to improve performance over selfish systems. The analysis considers performance of sum-rate and proportional fairness utility functions. This work assumes that the utility function is concave, which is not true for our application. Therefore, while this is the one of the most relevant to our work, it still does not study the scenario we do here.

Game Theory
An overview of the benefit of game theory for DSA is reviewed in [78]. This paper lists the several types of games that arise in different DSA scenarios (centralized vs. distributed and others) and performs a good overview of the benefits and needs of each.

In [79], a game theory approach derived with a genetic algorithm is used to distributively achieve the optimal average throughput to the binary symmetric channel using the physical model of interference. The paper notes that when the derived strategy is used against suboptimal strategies of other users, the derived strategy still achieves the best performance but it is not nearly as high as when all users use optimal strategies. Since this is a very special case scenario, it is not directly applicable to our work; however, the main strategies extracted from this analysis were segregation (avoiding bands with interference), robustness to exploitation (push me, I push you back), exploiting passive opponents, and randomize occasionally.

In [80], a noncooperative game is studied with a novel utility function which include SINR and cost for each channel. Three regions of interference exist (low or none, competitive, saturated). A Stackelberg game, a game that has a leader and followers instead of all users being equal, is modeled and the performance is examined briefly. Because this work does not consider any output metrics that we consider, it is not very relevant.

1.1.4 Cellular Networks

A special class of literature that may have applicable concepts to our topic includes the algorithms of cellular networks that optimize many parameters. Cellular research assumes the presence of many mobile users that connect to one of possibly many base stations. Cellular network research tends to divide optimization of user traffic into delay-tolerable and delay-intolerable [81]. An example of a delay intolerable application is voice traffic.

We found three main optimization functions that exist for cellular networks. The first one, used mainly for delay-intolerable applications, is a minimum signal to interference ratio, SIR, requirement [82]. In this optimization, as many users as possible are added to the system such that all users achieve the target SIR. Another power control algorithm commonly used for voice applications is to maximize the minimum SIR [83]. Instead of having a fixed SIR value like the minimum SIR optimization, the minimum SIR is variable and can be low with many users or high with few users. The final power control method optimization used in delay-tolerable conditions is that maximizing of the sum-rate or total capacity [84]. Since secondary DSA users are assumed to be delay tolerant, we will consider mostly the maximization of the sum-rate in this work.

1.2 Motivation

The spectrum scarcity problem provides the main motivation behind using Cognitive Radio and DSA. Several applications for Cognitive Radio and DSA have been introduced including military applications [85], emergency networks [86], industrial applications [87], and commercial wireless broadband [88]. In addition, several IEEE standards either have been or are in the process of being introduced that implement DSA or DSA technology including, SCC41 (also known as P1900), 802.11, 802.15, 802.16, 802.19, and 802.22 [89].

As mentioned previously, several implementation issues within the hierarchical model of DSA

still need generic solutions. As DSA has been suggested as a solution for inefficient spectrum usage of the primary users, we must ensure in the environments with several secondary users and scarce opportunistic bandwidth that secondary users efficiently use the leftover spectrum.

With the full usage of the ISM bands, it is likely that any DSA spectrum policy might have a high saturation of users. When there are more users desiring to use the channel than channels available, users will interfere with each other if left to greedy or random access schemes. If however, users cooperate or collaborate for frequency bands and transmit power, spectrum efficiency should be able to increase. However, this adds a layer of complexity in requiring all users to share the same collaboration protocol.

1.3 Contributions of this Thesis

We make the following set of contributions in this thesis:

- We characterize the benefit of using collaborative techniques over non-collaborative techniques (we particularly focus on a technique we call sensing) and show how this benefit is affected by different system parameters.
- We consider the impact of collaboration on the fairness of the distribution of capacity amongst the users. We also consider ways to improve fairness through collaboration.
- By determining what causes collaborative systems to achieve better performance than non-collaborative systems, we design a few non-collaborative approaches that are able to achieve closer performance to collaborative approaches than the original non-collaborative approach we consider.

Chapter 2

System Model

In this chapter, we introduce the system model that is the basis for all results in this report. The conceptual purpose of this chapter is two-fold: first, it explains the method we use to place transceiver nodes onto a plane (and the related assumptions) and second, it shows how we calculate individual user capacities given a specific simulation setup.

We begin by listing the set of assumptions that we make and defining the method used for the simulation setup. Next, we devote a section to how we calculate the channel gains for a system. We then introduce the different channel access techniques that we study and derive the sum-rate equations for each based upon the Shannon capacity of a link and treating all interference as Gaussian noise[1]. We conclude the chapter by introducing the different levels of collaboration that we will study in the following chapters.

2.1 Assumptions

We make the following set of assumptions for our system model:

- For each simulation, we assume that we have a set of paired nodes that have the ability to collaborate (which we call collaborative nodes). For a particular simulation, each pair of nodes includes a Tx (transmit) node and an Rx (receive) node, which together are called either a pair or a user.
- We model the spectrum by dividing it into a set of discrete uniform-sized orthogonal frequency bands which we call 'channels'. For the purpose of this work, we determine

[1] As is common in such studies, we use the sum-rate approach, because the optimal capacity for a set of paired nodes has only been determined for a limited number of specific cases.

the rate of the system in bits per second by assuming that each channel has a bandwidth of 1 Hertz. By dividing this rate by the number of channels, we determine the spectral efficiency. If the channel bandwidth of a system is not 1 Hertz, we can determine the actual rate of this system by multiplying the 1 Hertz rate by the actual channel bandwidth.

- For each simulation, we have a fixed number of channels, 'M', and a fixed number of users, 'N'.
- All Tx nodes have a maximum transmit power, P_{max}.
- Tx nodes can only control which channels they transmit on and how much power is transmitted on those channels; therefore, it is over transmit power and channel selection that all optimizations are performed.
- When users collaborate, they perfectly and instantaneously share channel, transmit power and rate information. The optimization is performed centrally and the optimal transmit power and channel selection factors are followed by all users.
- The path loss exponent, η, is equal to 3 in all simulations unless otherwise noted.
- Users are assumed to have a maximum transmit range d_{max} such that when the Tx node and Rx node are separated by d_{max} and the Tx node is transmitting P_{max}, the SNR at the Rx node neglecting fading and shadowing is equal to SNR_{min}. Tx and Rx nodes will never be separated by more than d_{max}. Since d_{max} is dimensionless, we remain general (i.e., $d_{max} = 1$ can represent one kilometer or one meter).
- If given the opportunity, a Tx node will transmit as much information as possible to its respective Rx node (thus we are assuming that the queue at the Tx node has infinite length).
- We use a physical model of interference (see [1]), and treat interference from other Tx nodes as Gaussian noise. We also assume the use of a Gaussian codebook which allows us to use the Shannon capacity formula per link as will be shown in Section 2.4.

2.2 Simulation Setup

In this section, we will describe step-by-step how Tx and Rx nodes are placed onto the plane for each simulation. We introduce two parameters, Multi-user Interference (MUI) and SNR_{min}, to scale the expected SIR and the expected SNR. Each iteration of the simulation uses the following steps (definitions which do not change are not necessarily repeated each step):

1. **We define d_{max}, select a value for SNR_{min} and the number of users N.** We set d_{max}, the maximum distance between a paired Tx and Rx node, to be 1 for all simulations. As mentioned previously, SNR_{min} is equal to the SNR at the receiver if the Tx node is transmitting maximum power, P_{max}, at a distance of d_{max} from the Rx node (neglecting fading or shadowing). For each simulation, we select a specific value for SNR_{min}. This will determine the range of the SNR distribution.

2. **We determine the noise power, ν.** We define the noise power ν according to the SNR_{min} parameter as shown in (2.1). Conceptually, we note that this is called SNR_{min} because, assuming transmit power of P_{max}, it is the minimum average (neglecting fading and shadowing) SNR possible since the Tx node must be within d_{max} of the Rx node.

$$\nu = \frac{P_{max}}{d_{max}^{\eta} \cdot SNR_{min}} \tag{2.1}$$

As SNR_{min} increases, the expected SNR increases, and as the SNR_{min} decreases, the expected SNR decreases. In Figure 2.1, we plot the cumulative density function (CDF) of the SNR for different values of SNR_{min} assuming a fixed $d_{max} = 1$ and $P_{max} = 1$. We can see that changing SNR_{min} results only in a shift of the CDF of the SNR (i.e., the curves are identical except for the SNR shift).

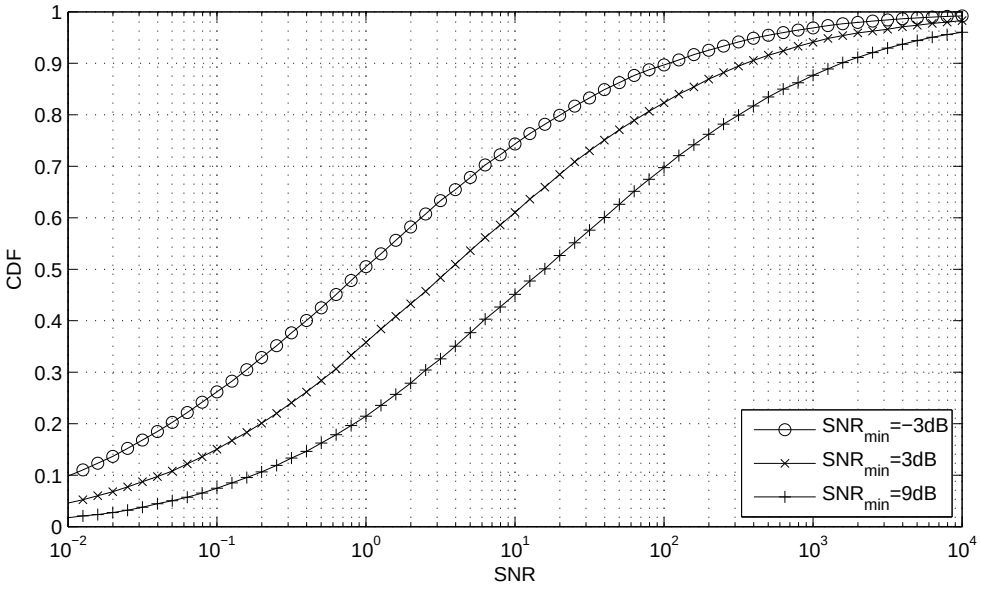

Figure 2.1: CDF of the SNR for different values of SNR_{min} ($N = 2$, $M = 1$, $d_{max} = 1$, $MUI = 2$ and $P_{max} = 1$)

3. **We select a value for the Multi-User Interference** factor (MUI). Conceptually, MUI represents the expected number of Rx nodes within a circle of radius d_{max} (assuming a constant density over an infinite plane). Increasing the MUI factor increases the number of expected interferers in a circle of radius d_{max}, increasing the interference and decreasing the SIR. Similarly, decreasing MUI decreases the interference and increases the SIR. In Figure 2.2, we plot the CDF of the SIR for different values of MUI for the case in which two users both transmit on a single channel. As expected, as we increase the MUI parameter the SIR values decrease.

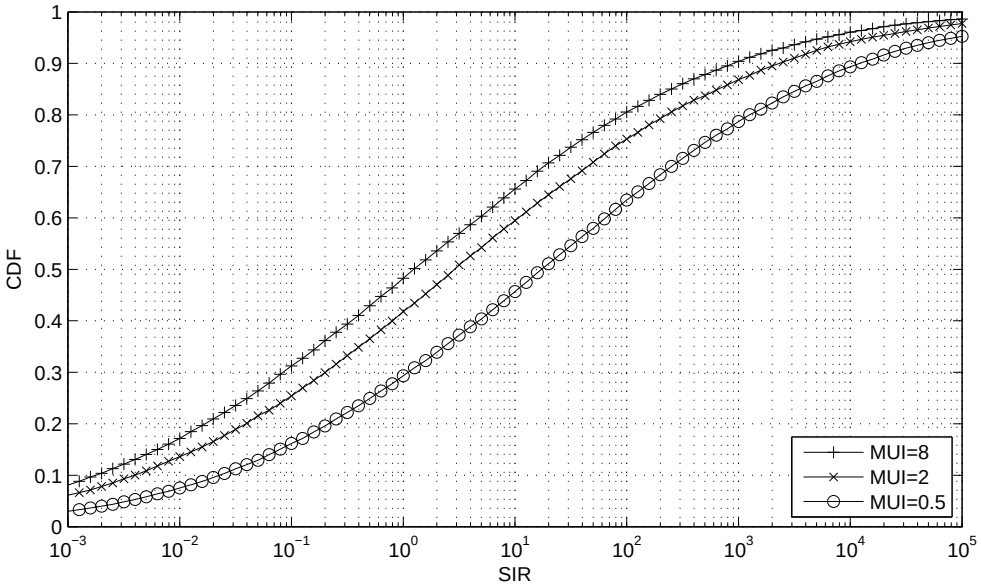

Figure 2.2: CDF of the SIR for different values of MUI ($N = 2$, $M = 1$, $d_{max} = 1$, $SNR_{min} = 3dB$ and $P_{max} = 1$)

We note that because MUI does not change d_{max}, P_{max}, or SNR_{min}, it does not change the distribution of SNR. Similarly, since the SNR_{min} factor does not affect MUI or d_{max}, it does not change the distribution of SIR. Therefore, these two factors remain independent and allow us to scale interference and noise independently. We can see this independence in the CDF of the SNR in Figure 2.4 for different values of the MUI parameter. We notice that the CDF of the SNR is independent of the MUI parameter (as evidenced by all CDF curves being identical). Similarly, we can see in Figure 2.3 that for different values of the SNR_{min} parameter the distribution of the SIR remains nearly the same. Thus the CDF of the SIR is independent of the SNR_{min} parameter.

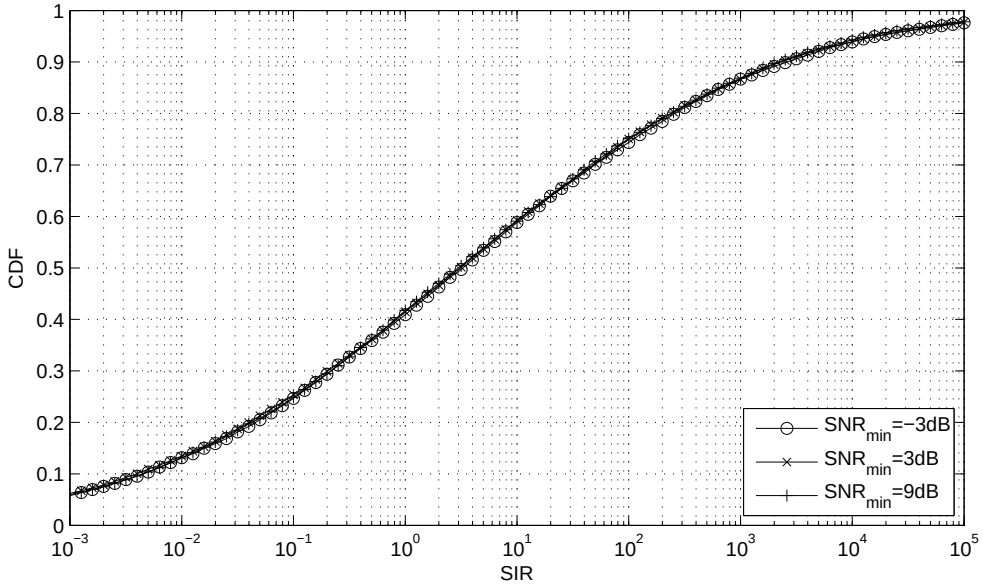

Figure 2.3: CDF of the SIR for different values of SNR_{min} ($N = 2$, $M = 1$, $d_{max} = 1$, $MUI = 2$ and $P_{max} = 1$)

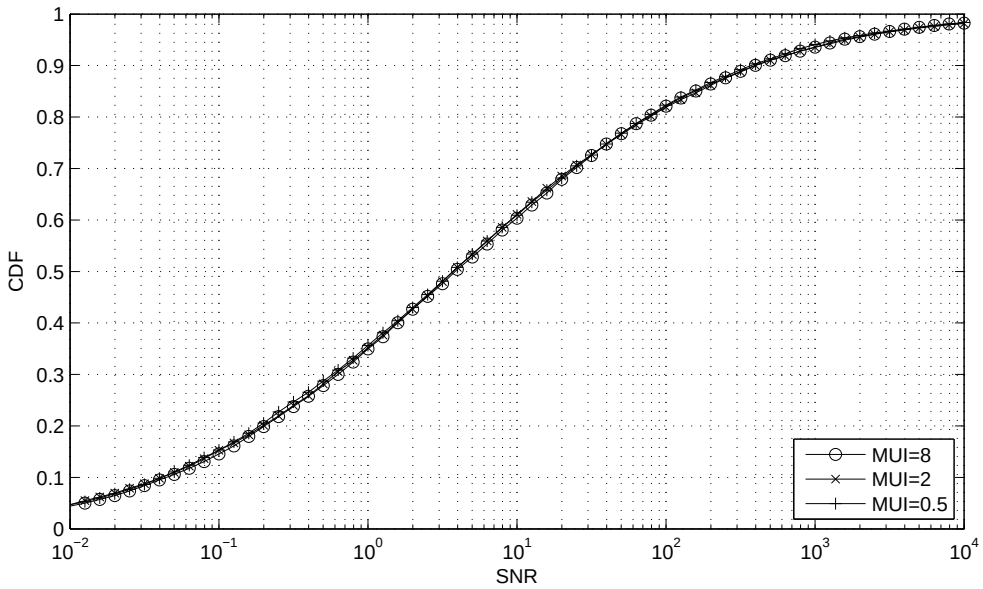

Figure 2.4: CDF of the SNR for different values of MUI ($N = 2$, $M = 1$, $d_{max} = 1$, $SNR_{min} = 3dB$ and $P_{max} = 1$)

4. **Using the MUI and d_{max} values, we determine the density that users are placed onto the plane.** MUI represents the expected number of Rx nodes within a circle of radius d_{max} assuming constant density across the plane. For a fixed MUI and d_{max} value, the density in terms of users per unit area, μ is given by (2.2).

$$\mu = \frac{MUI}{d_{max}^2 \cdot \pi} \tag{2.2}$$

5. **We create a square on a plane in which we place collaborative Rx nodes.** We note that we found no significant difference in performance between grouping users by the position of the Tx nodes or the position of the Rx nodes; therefore, without loss of generality, we group users by the position of the Rx node. We begin by creating a two dimensional plane upon which to place all Tx and Rx nodes. To place N users within a square with density of μ, the square must have an area of $\frac{N}{\mu}$. This means the length of a side of the square will be $\sqrt{\frac{N}{\mu}}$. All Rx nodes are placed uniformly inside this square (making the density equal μ).

6. **We place the collaborative Tx nodes.** We use d_{max} to place the Tx nodes. Each Tx node is placed uniformly within a circle of radius d_{max} centered at the respective Rx node. We note that, while all Rx nodes are required to be within the square, Tx nodes are only limited to being within d_{max} of the Rx node and may lie outside the square.

7. **If applicable, we add outside interference.** For some simulations, we add more Tx nodes that are not part of the collaboration which we will call interfering Tx nodes. We add interfering users by creating an outer square (which is concentric with the inner square and has a side length of 11 times the side length of the inner square) and defining a region inside the outer square but outside the inner square in which interfering nodes are placed. We found that making the side length of the outer square equal to 11 times the side length of the inner square approximates the impact of interference in an infinite distance in every direction (i.e., making the outside square any larger has little impact on the distribution of outside interference seen by collaborating users). We place interfering Rx nodes uniformly into this region, with a density, μ_{out}, related to the density of the collaborative users in the inner square, μ. To determine the value of μ_{out}, we define a loading factor $0 < \rho < 1$. The density for the outside interferers is $\mu_{out} = \mu \cdot \rho$. After placing Rx nodes with density μ_{out}, interfering Tx nodes are placed uniformly within a circle of radius d_{max} centered at the respective interfering Rx nodes and are assumed to select a channel randomly and transmit maximum power on that channel. We note that, similar to how the collaborative Tx nodes may lie outside the inner square, interfering Tx nodes may lie inside the inner square.

There are a few benefits to placing users on the plane the way we have. By using a constant d_{max}, we are able to scale this system accordingly to match any real system. By defining

MUI and SNR_{min} the way we have, we can keep either the expected SIR or expected SNR constant and determine the impact on the value of collaboration as we vary the other term (i.e., we can determine with a constant SNR distribution, the impact of increasing interference on the benefit of collaboration). Figure 2.5 shows an example layout of collaborative Tx and Rx nodes with no outside interference ($N = 7$ and $MUI = 2$).

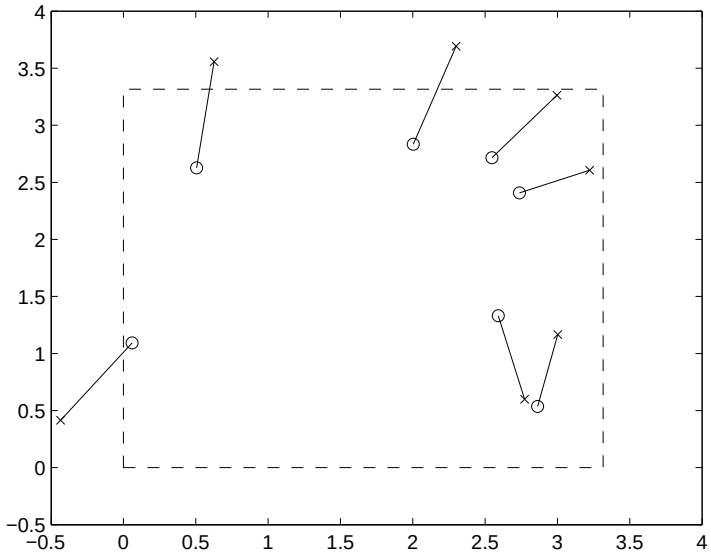

Figure 2.5: Layout of Tx and Rx Nodes. Tx Nodes are denoted by 'x' while Rx Nodes are denoted by 'o'. The dashed line indicates the edge of the collaborative square in which all Rx Nodes are placed.

2.3 Channel Gains

The distance from Tx node i to Rx node j is denoted as $d_{i,j}$. The channel gain from Tx node i to Rx node j on channel k is denoted $g_{i,j,k}$. Within the gain term, we have a path loss component, a multipath fading component (γ), and a shadowing component ($e^{\sigma \cdot N(0,1)}$ where $N(0,1)$ is a Gaussian random variable with mean of 0 and variance of 1). The gain term is derived as shown in (2.3).

$$g_{i,j,k} = \frac{1}{d_{i,j}^{\eta}} \cdot |\gamma|^2 \cdot e^{\sigma \cdot N(0,1)} \tag{2.3}$$

We assume that the small-scale fading component γ has an amplitude that is distributed according to a Rayleigh distribution and is independent across channels. The σ term in the

shadowing component determines the impact of shadowing ($\sigma = 0$ means no shadowing, and an increase in σ results in increasing shadowing effects). Shadowing may either be spatially independent or spatially correlated. In general, we assume that shadowing is perfectly correlated from frequency to frequency and independent from location to location.

When performing optimization in a real world scenario, fading may either be slow and predictable or fast and unpredictable. When fading is fast, channel gains are estimated by averaging the channel gain over a period of time which results in an estimate that includes path loss and shadowing only. Therefore, when we refer to optimizing in a fast fading environment, we optimize using a simplified gain term (i.e., neglecting short-term fading) and determine performance based upon the full gain term. For slow fading, we optimize and determine performance based upon the full gain term.

2.4 Channel Access Techniques

We examine the following three channel access techniques: single-channel (SC), multi-channel (MC) and spread spectrum (SS). In this section, we describe what we mean by each of these terms and how we determine an individual user capacity for each technique.

2.4.1 Terminology

This section describes the terminology used throughout the paper to represent capacities, powers, channel gains, etc. We denote the individual user capacity for user i across all channels as C_i and the capacity for user i on channel k as $C_{i,k}$. This means that,

$$C_i = \sum_{k=1}^{M} C_{i,k} \tag{2.4}$$

In a similar manner, the power and $SINR$ for user i on channel k are represented by $P_{i,k}$ and $SINR_{i,k}$.

2.4.2 SINR

Based upon our assumption that we treat interference as Gaussian noise and the assumption that we use a Gaussian codebook, our SINR for user i on channel k is given by (2.5), where $P_{i,k}^{rx}$ represents the received signal power and $I_{i,k}$ represents the total interference power.

$$SINR_{i,k} = \frac{P_{i,k}^{rx}}{I_{i,k} + \nu} = \frac{P_{i,k} \cdot g_{i,i,k}}{\sum_{j=1, j \neq i}^{M} P_{j,k} \cdot g_{j,i,k} + \nu} \tag{2.5}$$

2.4.3 Derivation of Capacity for Each Technique

We will derive capacity equations for each of the three techniques in this section. Because the capacity of the interference channel is unknown except for particular cases, we use Shannon capacity to determine the capacity of each link as shown in (2.6):

$$C_{i,j} = log_2\left(1 + SINR_{i,j}\right) \tag{2.6}$$

We note that this capacity equation is based upon our assumption that the channel bandwidth is 1 Hertz. We then define the "capacity" of the network as the sum-rate or sum-capacity of the network:

$$C = \sum_{i=1}^{M} C_i \tag{2.7}$$

Note that this capacity is in terms of bits per second per unit bandwidth.

Single-channel

With the single-channel access technique, all users are assumed to occupy a maximum of one channel. Accordingly, this means that each user i has only a single active channel l with a single data stream and the total capacity C_i is given by (2.8).

$$C_i = C_{i,l} = log_2\left(1 + SINR_{i,l}\right) \tag{2.8}$$

Multi-channel

In the multi-channel access technique, we assume that users can transmit on multiple channels simultaneously, sending a separate data stream on each channel. The total capacity of user i is given by (2.9).

$$C_i = \sum_{k=1}^{M} C_{i,k} = \sum_{k=1}^{M} log_2\left(1 + SINR_{i,k}\right) \tag{2.9}$$

Optimizing with the MC access technique is equivalent to optimizing with SC except that the user may use more than a single channel. Because the MC access technique encompasses the SC access technique but allows more freedom, we expect the performance of optimizing with MC to always be better than or equal to the performance of optimizing with SC.

Spread spectrum

For the spread spectrum technique, the user may transmit power on multiple channels, but only a single stream of data is sent (unlike MC where a stream is sent on each channel).

To allow for maximum flexibility, we assume frequency hopping techniques are employed since this allows any fraction of the power to be on any channel and does not require the channels to be continuous spectrum. (This is contrary to direct sequence spread spectrum where channels are continuous and the power is equally split to all channels).

The maximal SINR for spread spectrum techniques results from maximal ratio combining (MRC) across all channels. Using MRC combining causes the total $SINR_{i,SS}$ to equal the summation of the $SINR$ of the individual channels (this fact is derived in Appendix A). The $SINR_{i,SS}$ is shown in (2.10). Therefore, the capacity for the spread spectrum case is given by (2.11).

$$SINR_{i,SS} = \sum_{k=1}^{M} \frac{P_{i,k} \cdot g_{i,i,k}}{\sum_{j=1,j\neq i}^{N} P_{j,k} \cdot g_{j,i,k} + \nu} \tag{2.10}$$

$$C_i = log_2 \left(1 + SINR_{i,SS}\right) \tag{2.11}$$

Similar to the optimization with MC, optimizing with the SS access technique is equivalent to optimizing with SC except that the user may use more than a single channel. Because the SS access technique encompasses the SC access technique but allows more freedom, we expect the performance of optimizing with SS to always be better than or equal to the performance of optimizing with SC.

2.5 Collaboration Types

The three types of collaboration we study in this report include no collaboration (sensing), full collaboration, and partial collaboration. We will describe each here.

2.5.1 No Collaboration (aka Sensing)

The no collaboration case, which we will often refer to as *sensing*, corresponds to the case in which a user selects its power level and channel based upon local information to optimize an objective function. The user is limited to local information because no cooperation or collaboration occurs in this scenario. For our simulations, we assume that each user optimizes its own capacity selfishly based upon perfect $SINR$ estimation on each channel. Note that we could conceivably perform multiple rounds of selfish optimization/sensing in a round-robin fashion.

2.5.2 Full Collaboration

For full collaboration, a centralized optimization is performed over the power level and channel selection of all collaborative users simultaneously to optimize an objective function given the global channel gains and noise powers. The global information is assumed to be perfect and shared honestly. We will describe the objective function used in Section 4.1.

2.5.3 Partial Collaboration

The partial collaboration technique lies between the full and no collaboration cases. In partial collaboration, the users in the inner square are divided into groups where each group optimizes only among its members and groups are added sequentially. The optimization of each group is performed identically to the method described in section 2.5.2, but only the members of the group collaborate. Groups optimize given full information about the gains within the group and the sensed interference and noise power for each user on each channel due to other active groups.

This creates a suboptimal approach where all users might be willing but unable to fully collaborate due to spatial separation or using different protocols. Note that the case where the number of groups is one is equivalent to full collaboration and the case in which the number of groups is equal to the number of users is equivalent to sensing. The method we will use to place users into groups is the k-means (or Lloyd's) algorithm [90]. The k-means algorithm places users into groups according to geographic location (users which are close together tend to be placed into the same group). This algorithm is described in Table 2.1.

Table 2.1: K-means Algorithm

Algorithm 1 K-means Algorithm

Initialization:

1. Determine the number of users, N, and the number of groups, G. We define the position of user i to be H_i and the position of group j to be G_j. Each user i is associated with a group k_i. The distance between user i and its associated group is denoted by $|H_i - G_{k_i}|$. The number of activated groups at any point in time is denoted by n_G.
2. Place the center of the first group, G_1, at the centroid of all users N and associate all users i with group 1, $k_i = 1$ (the centroid location is determined by averaging over all users positions in each dimension).
3. If $G = 1$ then the algorithm is complete. Otherwise, find the user i^* that yields the largest value of $|H_{i^*} - G_1|$. Make $G_2 = H_{i^*}$. Set $n_G = 2$

Main Loop:

1. Re-associate each user i with the group j that corresponds to the minimum value of $|H_i - G_j|$ and assign $k_i = j$.
2. For each group j, re-calculate the centroid of all users associated with the group (i.e., G_j is the centroid of all users i^* where $k_{i^*} = j$).
3. If the user associations k_i and group centroid locations G_j have remained constant for two iterations, move on to Step 4. Otherwise, return to Step 1 of the Main Loop and repeat.
4. If $n_G < G$, find the user i^* that has the largest value of $|H_{i^*} - G_{k_{i^*}}|$ and set $G_{n_G+1} = H_{i^*}$ and $n_G = n_G + 1$ then return to Step 1 of the Main Loop. Otherwise, if $n_G = G$, the algorithm is complete with all users i are clustered according to their associated groups k_i with centers G_{k_i}.

2.6 Modulation and Coding Limitations

In real scenarios, systems have limitations given by the available modulation and coding techniques. We consider a few of these in this section.

2.6.1 Maximum Rate Limit

To mimic real world modulation/coding constraints, we have a maximum bits per second constraint that we implement for each user on each channel by modifying the individual $C_{i,j}$'s as shown in (2.12)

$$C_{i,j} = min\left(C_{i,j}, C_{max}\right) \tag{2.12}$$

In this equation, C_{max} is the maximum bits per second limit. For all simulations, $C_{max} = \infty$ unless otherwise stated.

2.6.2 Minimum Rate Limit

Similarly, we also have a minimum bits per second limit for each user on each channel. We implement this by modifying $C_{i,j}$ as shown in (2.13).

$$C_{i,j} = \left\{ \begin{array}{ll} 0 & C_{i,j} < C_{min} \\ C_{i,j} & otherwise \end{array} \right\} \tag{2.13}$$

In this equation, C_{min} is the minimum bits per second limit. For all simulations, $C_{min} = 0$ unless otherwise stated.

2.6.3 SINR Multiplier

The final modification to capacity that we consider is an SINR multiplier. The capacity calculations we use above assume that we can achieve the Shannon capacity on each link. Since this is not possible in practice, we multiply the SINR by a constant ξ, changing the capacity equation to (2.14) for SC and MC and (2.15) for SS where $0 < \xi \leq 1$. For all simulations, $\xi = 1$ unless otherwise stated.

$$C_{i,j} = log_2\left(1 + \xi \cdot SINR_{i,j}\right) \tag{2.14}$$

$$C_i = log_2\left(1 + \xi \cdot SINR_{i,SS}\right) \tag{2.15}$$

2.7 Metrics for Measuring Performance

In this section, we introduce metrics which allow us to compare the results of simulations with different input parameters. It has been extensively shown in the literature that there exists a tradeoff between capacity and fairness in systems like ours (e.g., [30]). Because of this, we introduce metrics that focus on capacity (sum-rate) and fairness (Jain Fairness). We also use outage capacity plots to illustrate the performance of individual users.

2.7.1 Sum-rate

A common metric used to measure capacity is sum-rate (or total capacity). As the true Shannon capacity of a distributed system (often termed the Gaussian Interference Channel) is not known, it is common to instead measure the sum of the capacities of individual links assuming Gaussian interference. This is written as

$$C = \sum_{i=1}^{N} C_i \tag{2.16}$$

To normalize the results to the number of users, we often use sum-rate per user (as given in (2.17)) instead of sum-rate. When comparing results, we often plot a cumulative density function (CDF) of the sum-rate or sum-rate per user to show how often the metric is below some value. In a few analyses, we use the term average sum-rate per user to denote the average of the sum-rate per user metric over several simulations.

$$\frac{\sum_{i=1}^{N} C_i}{N} \tag{2.17}$$

2.7.2 Jain Fairness

The Jain fairness metric measures fairness, i.e., the distribution of rates among users. A completely "fair" solution is one where all users obtain the same rate. Jain fairness is measured as

$$\frac{\left(\sum_{i=1}^{N} C_i\right)^2}{N \cdot \sum_{i=1}^{N} C_i^2} \tag{2.18}$$

When a single user has all of the capacity, the Jain fairness metric is equal to $\frac{1}{N}$ but when all users have the same capacity the Jain fairness metric is equal to 1. The optimal Jain fairness metric of 1 means that the capacity is fairly distributed amongst all users. The weakness of the Jain fairness metric is that some users may be unable to obtain high capacity, so increasing Jain fairness would require that the users with the best channels decrease their capacities instead of increasing the capacities of the users with the worst channels. As we

will see, improving fairness is typically at odds with maximizing sum-rate (capacity). When comparing simulation results, we often plot a CDF of the Jain fairness to denote how often the Jain fairness is below some value.

2.7.3 Outage Capacity Plots

Since sum-rate comparisons mask the performance of individual users, we also examine the distribution of the rate of individual users. When examining the tails of the CDF, we can determine the *outage capacity* of a user. As an example, please see the outage capacity plot in Figure 2.6. From the outage probability plot, we can see the performance of the users

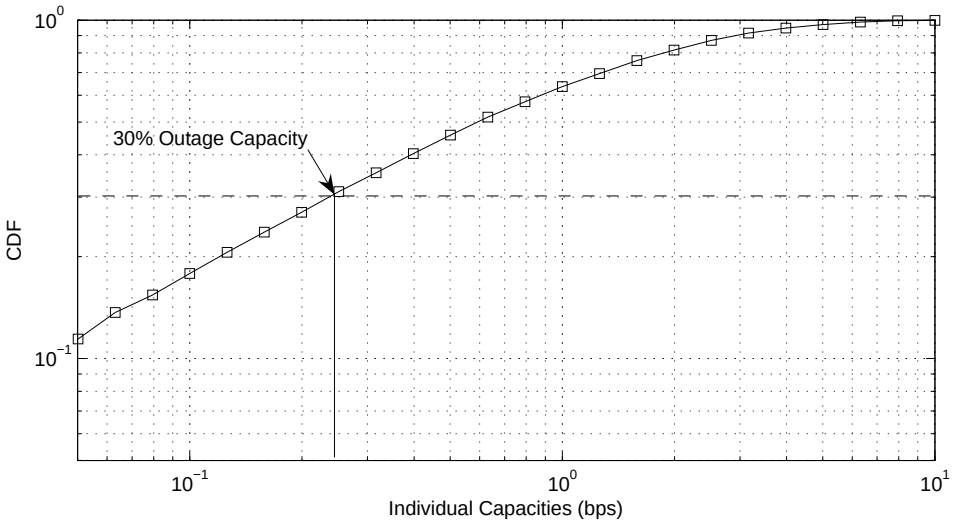

Figure 2.6: Sample Outage Probability Plot

with the worst capacities and the users with the best capacities simultaneously. We can also determine the probability for each technique that a given user will have a capacity over some specific value. As an example, in Figure 2.6 the 30% outage capacity is approximately 0.25 bps. Thus, the probability that a specific user achieves a capacity higher than 0.25 bps is 70% (or equivalently 70% of the users will get a capacity higher than 0.25 bps).

2.8 Summary

We began this chapter by stating the set of assumptions used in this work. We then described how each simulation begins by using a given set of parameters and following the procedure described in Section 2.2 to place Tx and Rx nodes onto a plane. After determining channel

gains using the equations in Section 2.3, we can calculate the C_i values as shown in Section 2.4 given the transmit power for each node on each channel ($P_{i,j}$). The level of collaboration determines whether the transmit power and channel selection factors are selected based upon local or global information. Finally, we introduce a set of metrics that we will use to compare simulation results in future chapters. In the next chapter, we present simulation results for the non-collaborative approach which we call sensing.

Chapter 3

Modeling Non-Collaboration: Sensing

In the previous chapter, we introduced the system model and the metrics we use to compare the performance of simulations with different sets of parameters. In this chapter, we study the sensing approach in more detail. In the first section, we explain the impact of the number of users and number of channels on the capacity of sensing. We then study the impact of allowing users to use multiple channels in sensing by introducing a multi-channel sensing approach. In the next section, we study the impact of allowing users optimize their transmit power and channel selection multiple times in a round-robin fashion (we call this multiple rounds of sensing). We then consider the fairness of sensing approaches. We consider the impact of our SNR_{min} and MUI parameters (noise and interference) on the performance of sensing. Finally, we consider an approach where users are required to achieve a target rate to transmit on a channel.

3.1 Multiple Users / Multiple Channels

In this section, we examine the impact of the number of users and channels on the average sum-rate per user of sensing. For the case of $SNR_{min} = 3dB$ and $MUI = 5$, we have plotted the average sum-rate per user for different numbers of users and channels in Figure 3.1. As we increase the number of channels, the average sum-rate per user increases. The increase in average sum-rate per user is due to a benefit of the interference power being spread among more channels and a gain in multi-channel diversity. For example, for the case in which 6 users share a single channel, the average sum-rate per user is very small (about 0.9 bps) due to large cochannel interference; however, when an additional channel is added, the number of users active on each channel is cut in half causing significant decrease in interference. Furthermore, additional channels give a multi-channel diversity gain by giving users more channels to select from. From the figure, we also notice that, as the number of users increases, the sum-rate per user decreases. This decrease in the capacity per user as the number of

users increases is a commonly expected result from information theory (i.e., [1] although our system model is different). This decrease is due to the increasing interference per channel as the number of users increases. We note that, as the number of users increases, the expected interference power seen at a user will approach a constant (assuming a path loss factor larger than 2) causing the average sum-rate per user to approach a constant as the number of users goes toward infinity.

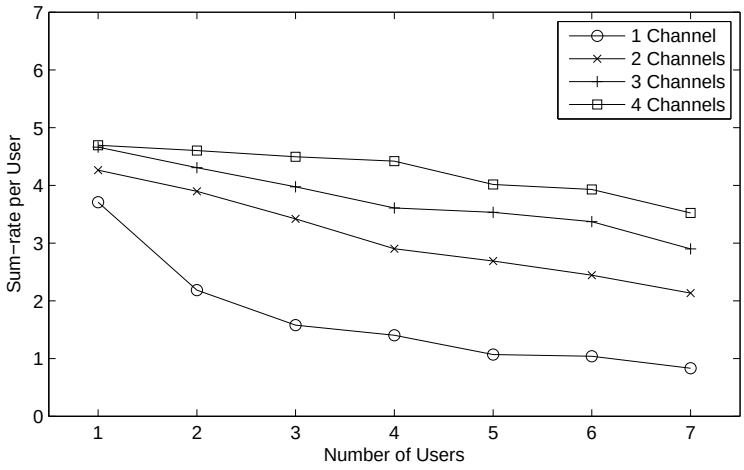

Figure 3.1: Sum-rate per User for Sensing with $SNR_{min} = 3dB$ and $MUI = 5$

We note that a theoretical capacity expectation has been derived by Gupta and Kumar in [1] for a system model similar to our own. The main differences between our system model and that of Gupta and Kumar is that as the number of users increases, our user density remains constant (for Gupta and Kumar it increases), and we place users into a square while Gupta and Kumar place users into a circle. From the definition of MUI, we note that our model should nearly match that of Gupta and Kumar when the number of users is equal to the MUI (by nearly, we mean that the only difference remaining is that we use a square while Gupta and Kumar use a circle). When $N < MUI$, the interference for our model should be higher than the interference for the model of Gupta and Kumar, so our average sum-rate per user for this scenario should be lower than the theoretical value. When $N > MUI$, the interference for our model should be lower than the interference for the model of Gupta and Kumar, so our average sum-rate per user for this scenario should be higher than the theoretical value.

We plot the comparison of the sensing technique with a single channel to the theoretical bound of Gupta and Kumar in Figure 3.2. We notice that the performance of sensing decays much faster than the theoretical bound. This is due to the fact that the sensing approach is not optimal, while the bound is intended to predict the optimal result. Therefore, we do not expect the result of sensing to follow the theoretical bound introduced by Gupta

and Kumar; we will consider this bound again later when considering the results of full collaboration (which should yield optimal results and achieve close to the bound).

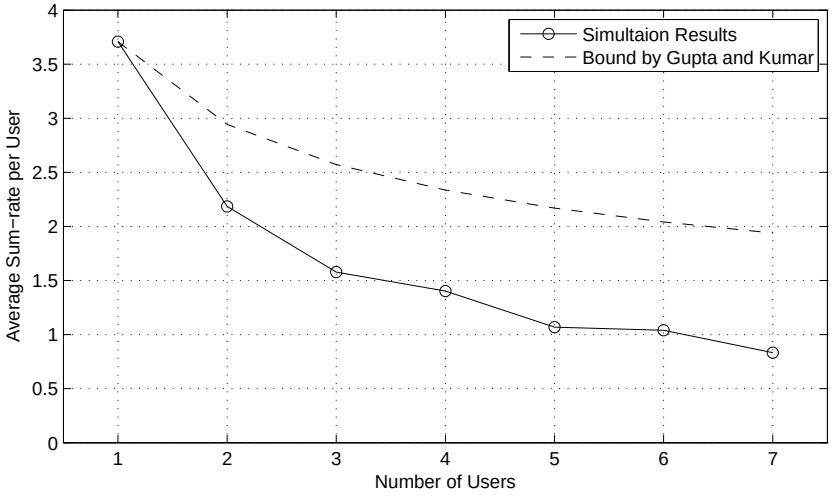

Figure 3.2: Comparison of Simulation Results to Gupta/Kumar bound ($SNR_{min} = 3dB$, $MUI = 5$ and $M = 1$)

In Figure 3.3, we plot the CDF of the sum-rate per user for 6 users with 1 and 4 channels. When we have 4 channels, the sum-rate per user values tend to be higher due to the ability of the users to avoid causing harmful interference with each other. When we have a single channel, all users transmit full power on that channel causing significant harmful interference to each other. We can see in Figure 3.4 (the CDF of the Jain fairness) that the fairness is significantly higher with four channels than it is with a single channel. This is due to the fact that when there is a single channel fewer users share high performance whereas when four channels are present many users achieve good performance. We can see this fact in the CDF of the individual capacities in Figure 3.5. The steeper slope of the CDF for four channels suggests that more users achieve similar high capacities which makes fairness higher (for example, 22% of users achieve higher than 1 bps when a single channel is present but 68% of users achieve higher than 1 bps when four channels are present).

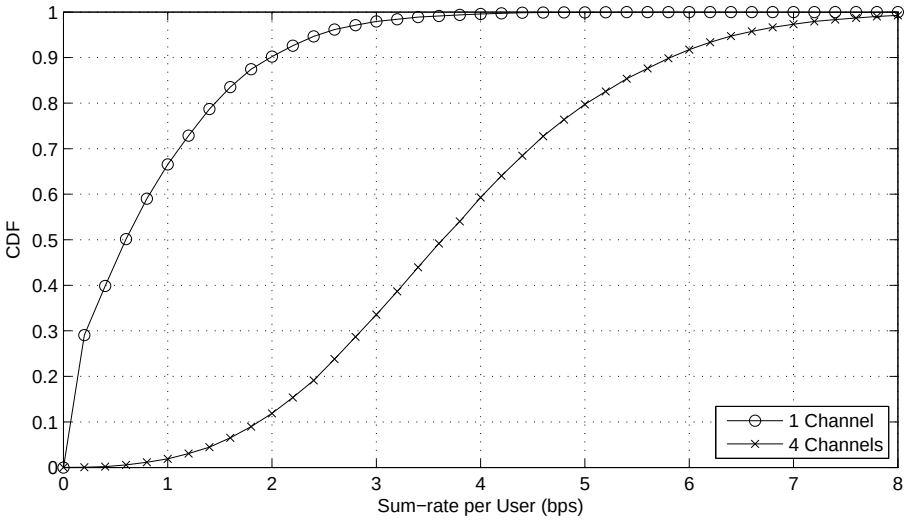

Figure 3.3: CDF of Sum-rate per User for Sensing ($SNR_{min} = 3dB$, $MUI = 5$ and $N = 6$)

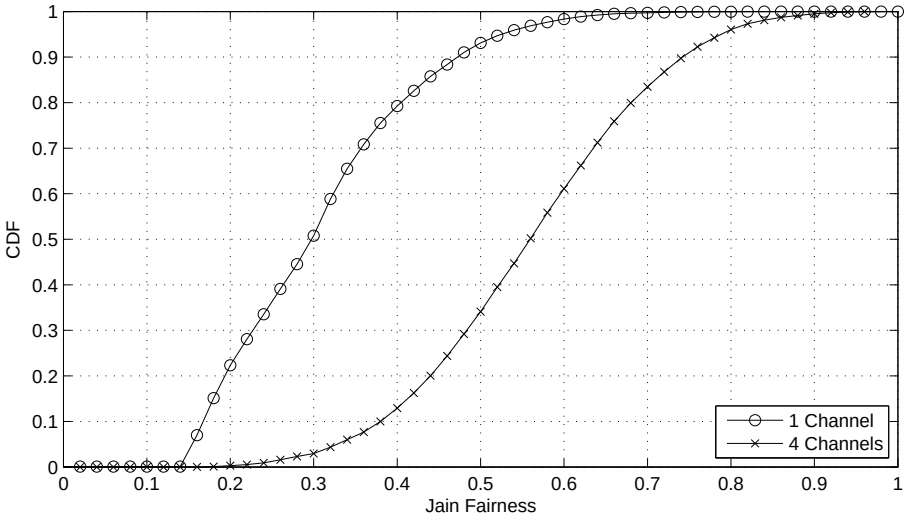

Figure 3.4: CDF of Jain Fairness for Sensing ($SNR_{min} = 3dB$, $MUI = 5$ and $N = 6$)

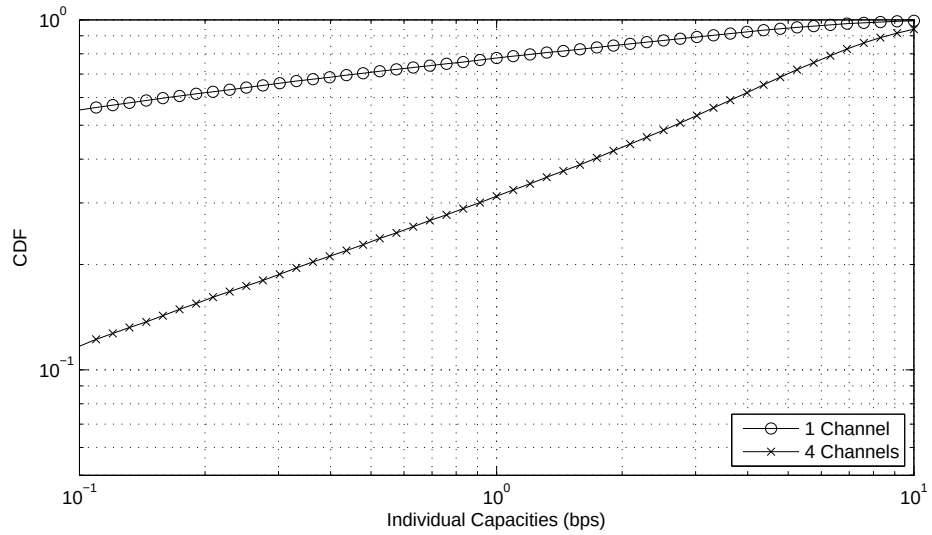

Figure 3.5: CDF of Individual Capacities for Sensing ($SNR_{min} = 3dB$, $MUI = 5$ and $N = 6$)

3.2 Multi-channel sensing

We have assumed up until now that for sensing a user senses each channel and transmits on the channel that gives it the best performance. However, we now consider what benefit we could gain by allowing a user to sense the available channels and determine the optimal power allocation amongst all available channels using a multi-channel technique with an independent data stream on each channel. Similar to our previous sensing approach, we assume that users optimize their multi-channel power allocation selfishly and are added sequentially. When using multi-channel sensing, we maintain that the maximum transmit power a user can have is P_{max} across all channels. We determine the optimal power allocation by using an optimization routine to maximize user capacity. We note that we could have performed this optimization by using waterfilling; however, we have found that our results with the optimization routine are nearly identical to the result using waterfilling.

To compare the performance of the multi-channel and single-channel sensing approaches, we plotted that benefit as a ratio of the average sum-rate per user of multi-channel sensing to the average sum-rate per user of single-channel sensing in Figure 3.6 for a varying number of users and channels. When the number of channels is larger than the number of users, we notice that a considerable benefit exists for using multi-channel sensing compared to single-channel sensing (typically over 50%). This benefit is mainly due to the inability of single-channel sensing to fully utilize the available channels (it can only utilize as many channels as there are users); multi-channel sensing is able to utilize these channels. As we increase the number of users, this benefit tends to decrease quickly. When many users are trying to share a few channels, the system tends to be interference limited causing multi-channel and single-channel approaches to both perform almost equally poorly.

3.3 Multiple Rounds of Sensing

In this section, we explore the benefit of using multiple rounds of sensing. By multiple rounds of sensing, we mean that, once all users have sequentially chosen a channel (as done in our traditional definition of sensing), we allow users to sequentially re-optimize their performance in an attempt to improve performance. For example, in a two user system, if user 1 selects a channel and transmits on that channel and then user 2 selects a channel and transmits on that channel, a second round of sensing would involve user 1 re-evaluating the channels (now that user 2 is transmitting) to try to determine if better performance is possible on another channel followed by user 2 following the same re-evaluation process in search of better performance (this would be 2 rounds of sensing). We note that, while the user that re-evaluates the channel most recently may have better performance on average than the other users, the users are randomly ordered, so this should not impact the results of our simulations (i.e., rerunning the simulation with a different sensing order should yield an equivalent sum-rate

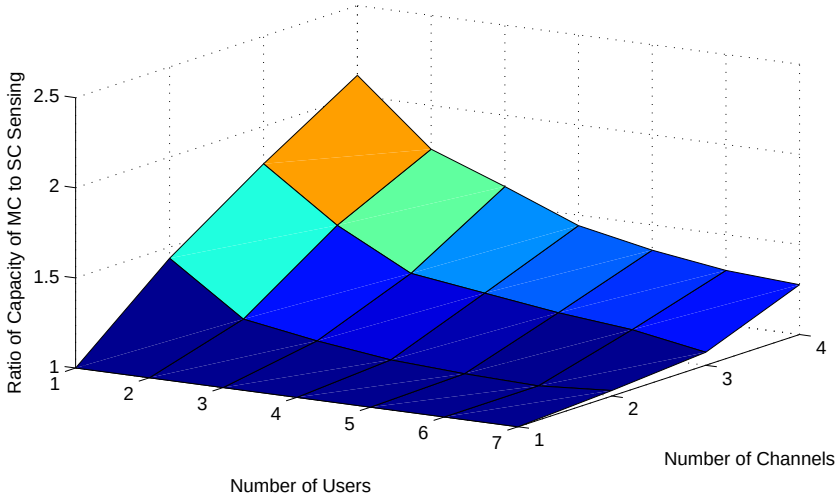

Figure 3.6: Ratio of Sum-rate of MC Sensing to the Sum-rate of SC Sensing with $SNR_{min} = 3dB$ and $MUI = 5$

per user on average). We note that the only benefit that can be achieved through multiple rounds of sensing is better channel selection (because no power control is used in sensing).

For the case of $M = 2$, $SNR_{min} = 3dB$ and $MUI = 5$, we plot the CDF of sum-rate per user for 1 round, 2 rounds, and 10 rounds of SC sensing for $N = 4$ in Figure 3.7 and $N = 7$ in Figure 3.8. We immediately notice that, while there is some benefit to multiple rounds of sensing, the benefit is quite small. About half of the benefit is achieved by using two rounds of sensing. The median sum-rate per user for $N = 4$ increases from about 2.5 to about 2.6 (about 4% increase) while the median sum-rate per user for $N = 7$ increases from about 1.9 to about 2.0 (about 5% increase). These gains are relatively small because a single round of sensing achieves nearly the best channel selection possible through sensing. We plot the ratio of the average sum-rate per user for SC sensing with 10 rounds to the average sum-rate per user for SC sensing with 1 round of sensing in Figure 3.9 for $SNR_{min} = 3dB$, $MUI = 5$ and a varying number of users and channels. From this plot, we can see that the maximum benefit of multiple rounds of sensing is a little more than 5%. The benefit tends to remain relatively flat for most numbers of users and channels while $N > 1$ and $M > 1$. Therefore, considering that extra sensing will require more overhead with little benefit, there is almost no benefit to performing additional sensing beyond the spectrum sensing requirement for DSA systems.

For the case of $M = 2$, $SNR_{min} = 3dB$ and $MUI = 5$, we plot the CDF of the sum-rate per user for 1 round, 2 rounds, and 10 rounds of multi-channel sensing with $N = 4$ in Figure 3.10 and $N = 7$ in Figure 3.11. From these figures, we can see that the benefit of multiple rounds of sensing is higher for multi-channel sensing than it was for single-channel sensing,

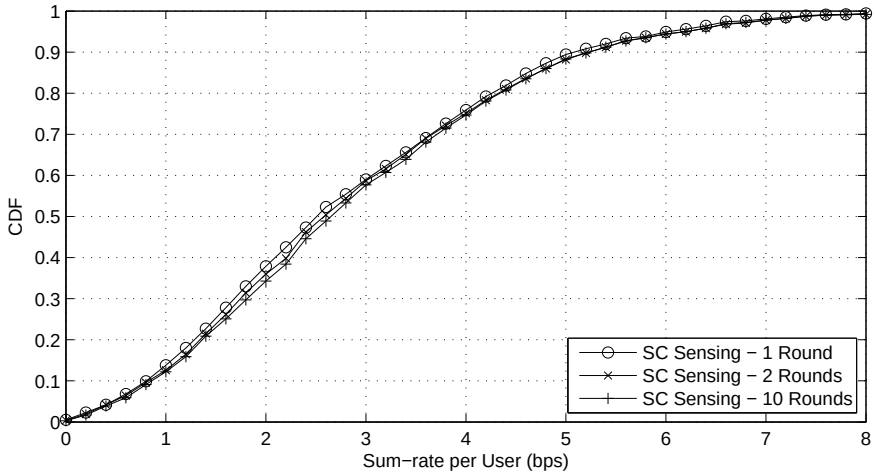

Figure 3.7: CDF of Sum-rate per User of SC Sensing with 4 Users and 2 Channels ($SNR_{min} = 3dB$ and $MUI = 5$)

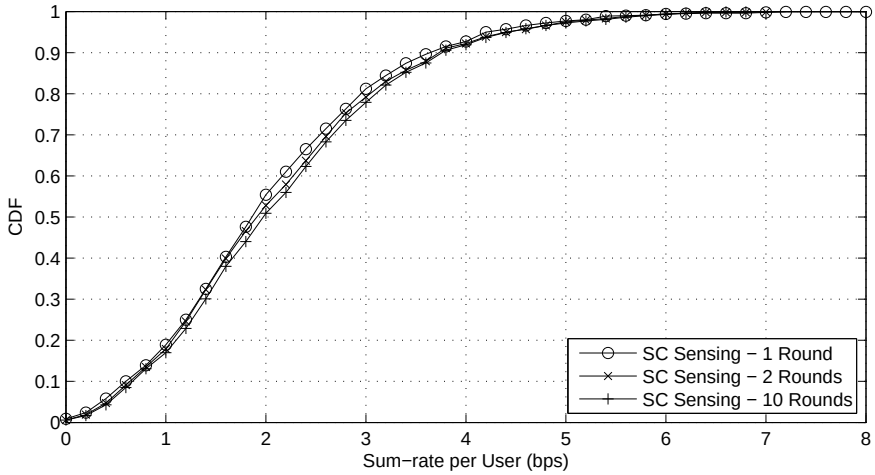

Figure 3.8: CDF of Sum-rate per User of SC Sensing with 7 Users and 2 Channels ($SNR_{min} = 3dB$ and $MUI = 5$)

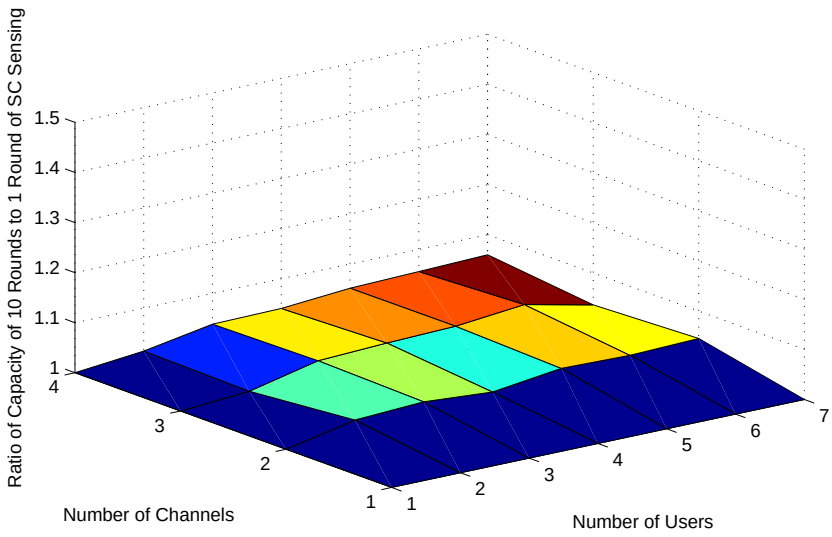

Figure 3.9: Ratio of Sum-rate of SC Sensing with 10 Rounds to the Sum-rate of SC Sensing with 1 Rounds ($SNR_{min} = 3dB$ and $MUI = 5$)

but the gains of multiple rounds are still small. Most of the gains achieved by multiple rounds of sensing with multi-channel sensing can be achieved with only 2 rounds of sensing. For $N = 4$, the median sum-rate per user increases from about 2.4 bps to about 2.6 bps (a gain of about 8%), and for $N = 7$, the median sum-rate per user increases from about 1.8 bps to about 2.0 bps (a gain of about 11%). We plot the ratio of the average sum-rate per user of multi-channel sensing with 10 rounds to the average sum-rate per user of multi-channel sensing with 1 round in Figure 3.12 for $SNR_{min} = 3dB$, $MUI = 5$ and a various number of users and channels. We can see that the maximum benefit of using multiple rounds of sensing is a little more than 10%. The benefit of multiple rounds with multi-channel sensing tends to increase with the number of users and with the number of channels. This means that a single round of multi-channel sensing becomes less likely to be the best that sensing can perform as the complexity of the system increases (complexity increases because more users causes more interaction among users and more channels allows more options for each individual user).

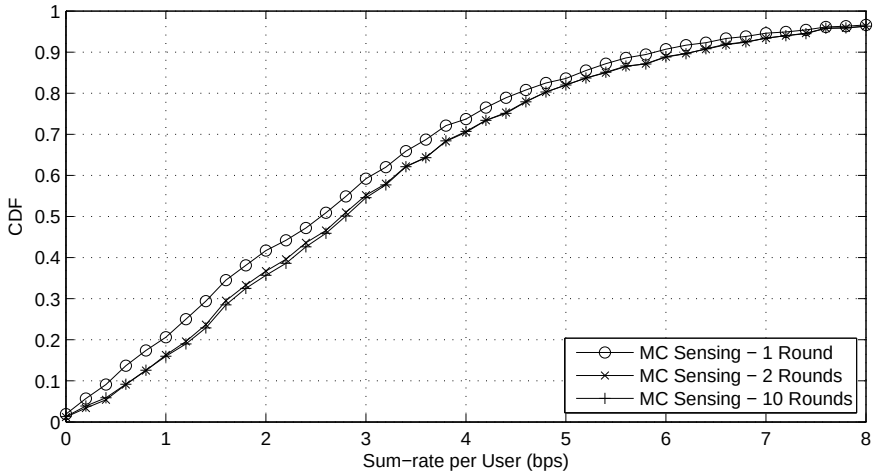

Figure 3.10: CDF of Sum-rate per User of MC Sensing with 4 Users and 2 Channels ($SNR_{min} = 3dB$ and $MUI = 5$)

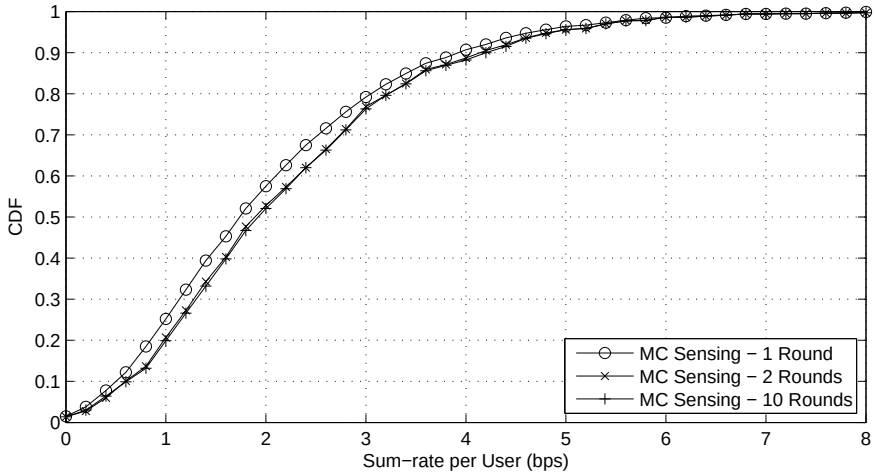

Figure 3.11: CDF of Sum-rate per User of MC Sensing with 7 Users and 2 Channels ($SNR_{min} = 3dB$ and $MUI = 5$)

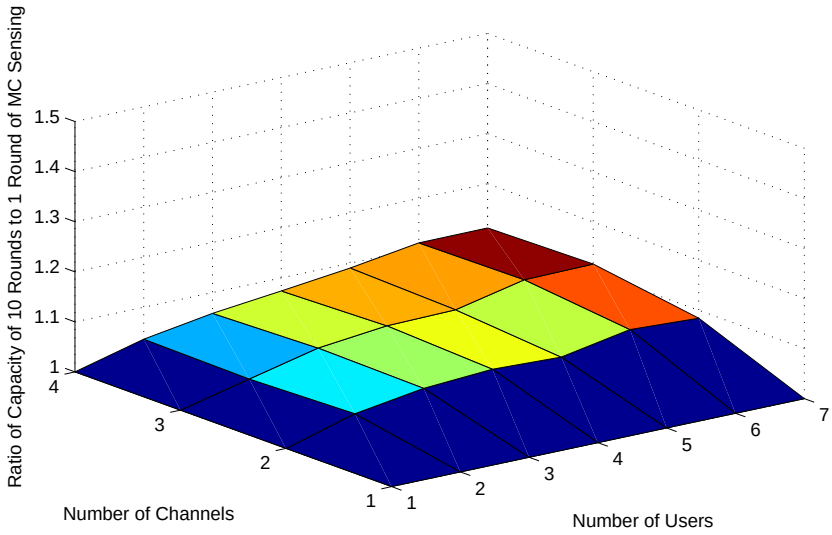

Figure 3.12: Ratio of Sum-rate of MC Sensing with 10 Rounds to the Sum-rate of MC Sensing with 1 Rounds ($SNR_{min} = 3dB$ and $MUI = 5$)

3.4 Fairness of Sensing

We now consider the fairness of the sensing approaches. In a fair system, all users tend to achieve similar throughput. In Figure 3.13, we show the CDF of the sum-rate per user to illustrate the fairness of sensing. For the best 2 user and best 4 user case, we simply sum the rates of the two or four users with the best rate and divide this sum-rate by the total number of users in the system. By comparing the CDF curves in which two users, four users, and all user are considered we can determine how much impact individual users have on the sum-rate per user of the system. The average sum-rate when we only consider the best two users is 1.62 bps, the average sum-rate when we only consider the best four users is 2.11 bps, and the average sum-rate per user when we consider all users is 2.22 bps. Therefore, the best two users achieve around 75% of the total rate of the system, the middle two users achieve around 20% of the rate of the system, and the worst users achieve around 5% of the rate of the system. From this comparison, we can see that this scenario has poor fairness (the best users achieve a significant portion of the capacity and the worst users get almost no capacity).

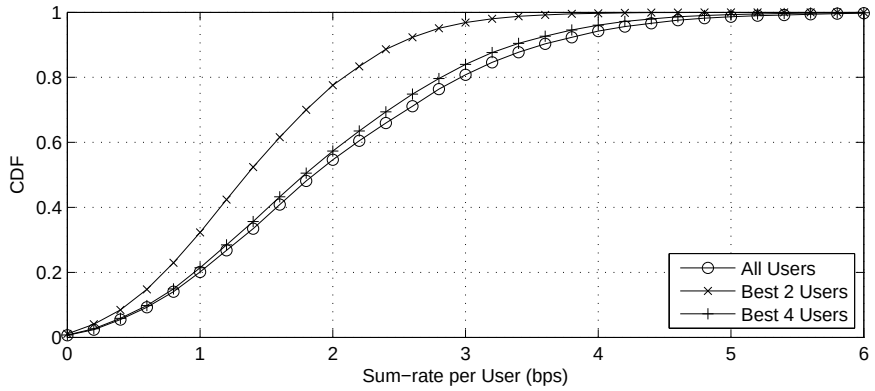

Figure 3.13: CDF of Sum-rate per User for 7 Users and 2 Channels ($SNR_{min} = 3dB$ and $MUI = 5$). For Best 2 Users and Best 4 Users, the Sum-rate per User is Calculated by Summing the Rates of the Best 2 or 4 Users and Dividing by 7.

In Figure 3.14, we can see the average Jain fairness for varying numbers of users and channels. We can see that as the number of users increases, the fairness tends to decrease. As we increase the number of channels, the fairness tends to increase. These two results mean that we should expect that as the interference increases the fairness will decrease, but as the interference decreases the fairness will increase.

In Figures 3.15 and 3.16, we compare the CDF of the Jain fairness for multi-channel and single channel approaches for 7 users and 4 users respectively. We can see for both cases that multi-channel sensing has worse fairness than single-channel sensing. This is due to the

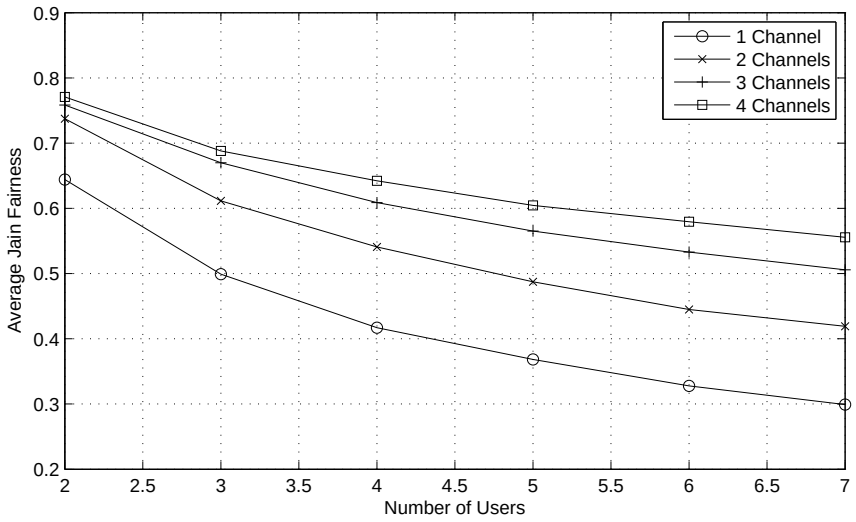

Figure 3.14: Average Sum-rate per User for Varying Numbers of Users and Channels ($SNR_{min} = 3dB$ and $MUI = 5$)

fact that users with better SNR values achieve higher gains from multi-channel approaches (we will show this in Section 5.7). This means that the best users can achieve significantly better capacities with more channels while the capacities of the worst users can only improve slightly. Because users tend to spread their power across the available channels with multi-channel sensing, there is a loss in multi-channel diversity (if user a causes user b significant interference, user b may no longer be able to avoid user a with user a transmitting on multiple channels). In Figure 3.17, we plot the outage probability to view the individual user capacities for 4 users and 2 channels. We can see that the capacities of the best users are higher using multi-channel technique (the 90% outage capacity of MC sensing is about 8 bps while 90% outage capacity of SC sensing is about 6.5 bps), and the capacities of the worse users are lower using the multi-channel technique (the 50% outage capacity of MC sensing is about 0.7 bps while the 50% outage capacity of SC sensing is about 1.2 bps).

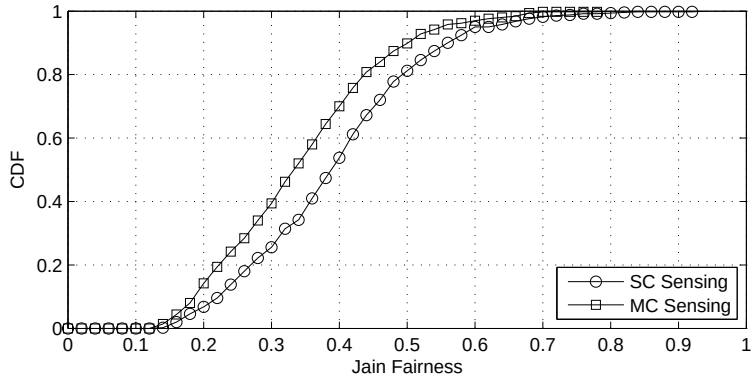

Figure 3.15: CDF of Jain Fairness of MC and SC Sensing for 7 Users and 2 Channels ($SNR_{min} = 3dB$ and $MUI = 5$)

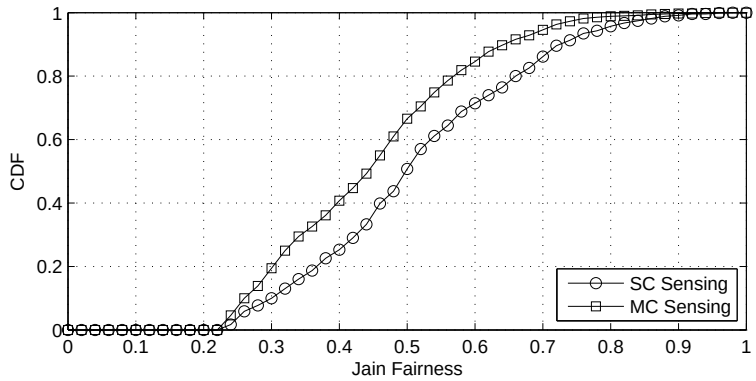

Figure 3.16: CDF of Jain Fairness of MC and SC Sensing for 4 Users and 2 Channels ($SNR_{min} = 3dB$ and $MUI = 5$)

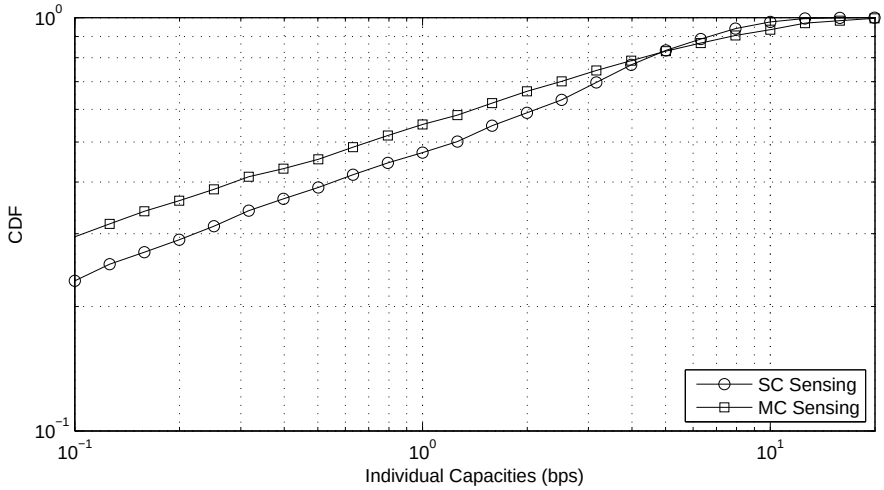

Figure 3.17: CDF of Individual User Capacity of MC and SC Sensing for 4 Users and 2 Channels ($SNR_{min} = 3dB$ and $MUI = 5$)

In Figure 3.18, we plot the CDF of Jain fairness of SC sensing for 7 users and 2 channels for 1 and 10 rounds. We can see that multiple rounds of sensing tends to increase fairness (remember that increasing the number of rounds also increases sum-rate per user). We can see why using multiple rounds of sensing tends to increase capacity and fairness by looking at the outage capacity plot for the same parameters in Figure 3.19. We can see that using multiple rounds has no impact on the capacities of the users with the best capacities, but it increases the capacities of the users with the worse capacities. We can see a similar impact by considering the CDF of Jain fairness and outage probability plot for MC Sensing with 7 users and 2 channels in Figures 3.20 and 3.21. (MC sensing also has a simultaneous small increase in sum-rate per user and small increase in Jain fairness.) From the outage capacity plot, we can see that this benefit is also due to an increase in the capacities of the users with the worst capacities.

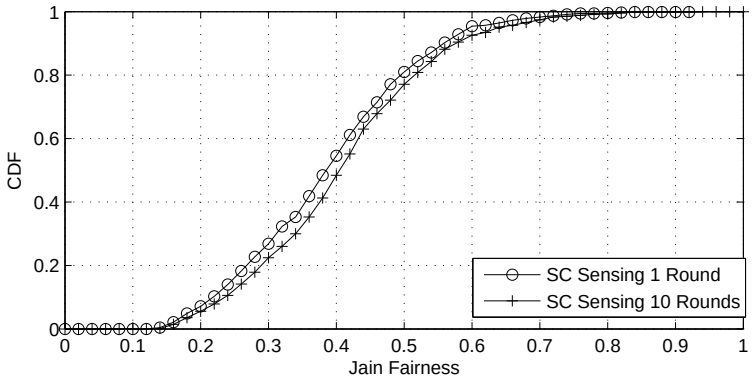

Figure 3.18: CDF of Jain Fairness of SC Sensing with Multiple Rounds for 7 Users and 2 Channels ($SNR_{min} = 3dB$ and $MUI = 5$)

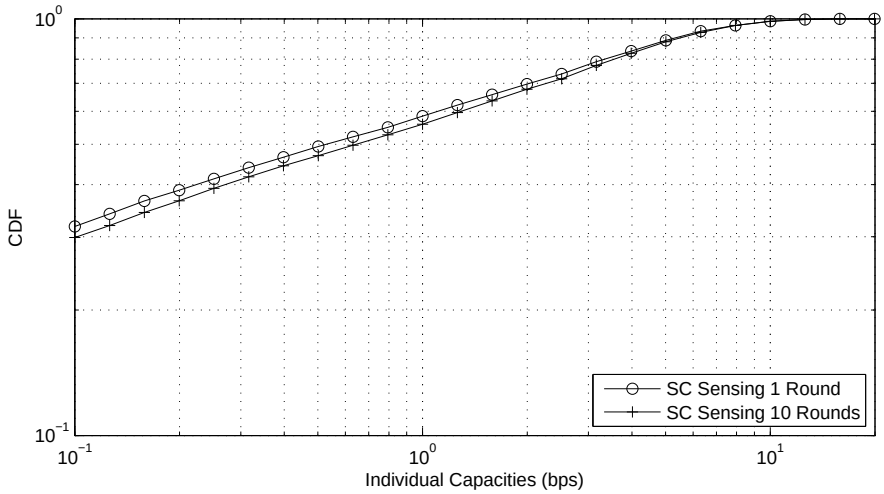

Figure 3.19: CDF of Individual User Capacity of SC Sensing with Multiple Rounds for 7 Users and 2 Channels ($SNR_{min} = 3dB$ and $MUI = 5$)

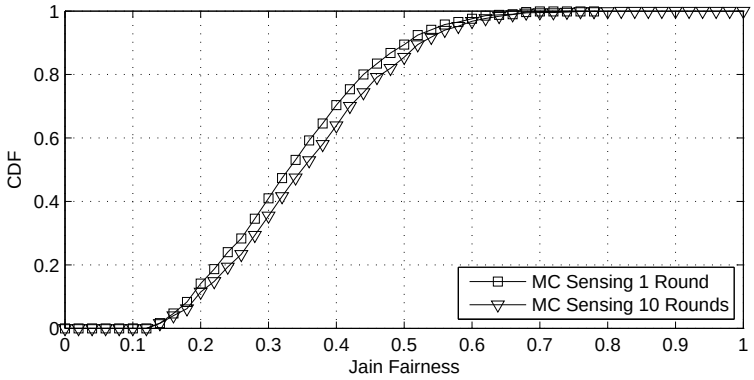

Figure 3.20: CDF of Jain Fairness of MC Sensing with Multiple Rounds for 7 Users and 2 Channels ($SNR_{min} = 3dB$ and $MUI = 5$)

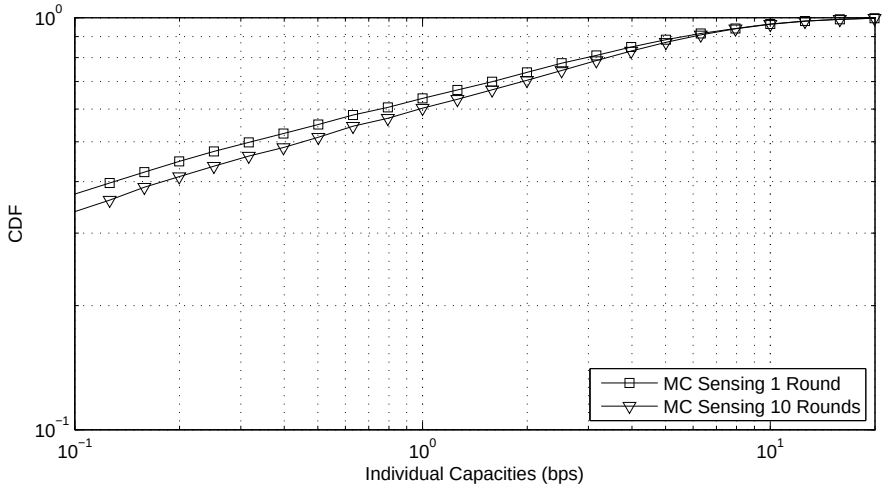

Figure 3.21: CDF of Individual User Capacity of MC Sensing with Multiple Rounds for 7 Users and 2 Channels ($SNR_{min} = 3dB$ and $MUI = 5$)

3.5 Impact of Interference and Noise

In this section, we determine the impact of changing the MUI and SNR_{min} factors (effectively changing the expected interference power and expected noise power). In Figure 3.22, we plot the average sum-rate per user for $N = 6$ and $M = 2$ with varying levels of interference and noise. As we increase the expected interference power (i.e., increasing the density) by increasing the MUI, we can see that the performance degrades as expected. The largest decrease occurs as we increase the interference from almost none ($MUI = 0.1$) to a small amount ($MUI = 1$). Similarly, as we increase the expected noise power (or decrease the SNR_{min}), the performance is also degraded, but this degradation seems to follow a nearly linear decrease for a fixed MUI. For example, for $MUI = 1$, the average sum-rate per user for $SNR_{min} = -3dB$ is 3.3; however, as we increase the SNR_{min} by $6dB$ the average sum-rate per user increases to 4.1 and a further increase of $6dB$ causes the average sum-rate per user to increase to 4.9 (making an increase in $6dB$ relate to an average sum-rate per user increase of 0.8 bps for $MUI = 1$).

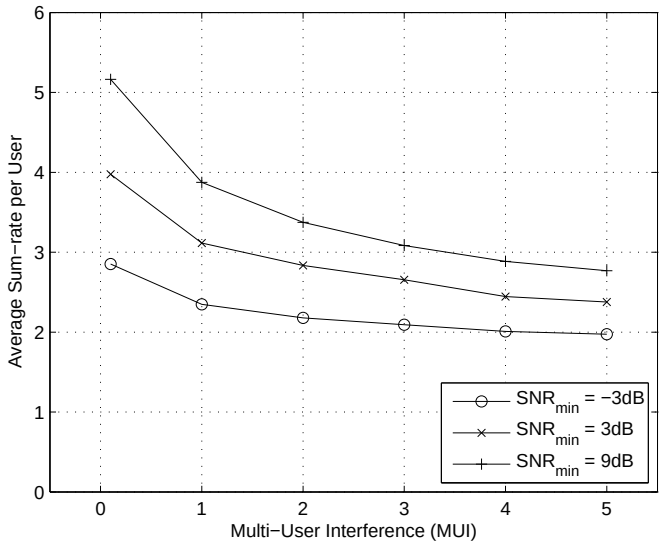

Figure 3.22: Impact of Interference and Noise on the Average Sum-rate per User of Sensing ($N = 6$ and $M = 2$)

3.6 Meeting a Target Rate

In this section, we consider the performance of a target-based system. By saying a target-based system, we specifically consider the situation in which all users are only allowed to achieve a fixed capacity C_{target}. In other words, $C_{min} = C_{target}$ and $C_{max} = C_{target}$. In this type of system, instead of all users transmitting simultaneously, users only transmit if they can achieve the C_{target}. While it may seem that in this system allowing users to transmit any power P such that $0 \leq P \leq P_{max}$ would help to minimize the interference of the system (by users only transmitting the power necessary to achieve C_{target}), this leads to instability. For example, if user a is transmitting on a channel with the minimal power necessary to achieve C_{target} and user b begins to transmit on the same channel with the minimal power necessary, user a will no longer achieve C_{target}. Therefore, user a will have to increase its transmit power slightly to achieve C_{target} again, but this will cause user b to no longer achieve C_{target}. Because this iterative competition of power levels is not guaranteed to converge, we instead assume that each user can either transmit no power or the maximum power ($P = 0$ or $P = P_{max}$). In this system, if user a is transmitting on a channel with the maximum power and user b begins to transmit on the same channel with the maximum power, user a will either be able to achieve C_{target} or fail to achieve C_{target}. If user a is unable to achieve C_{target}, it is simply cut off from using the channel. This yields to a much simpler algorithm. Namely, users are sequentially added to channels if possible, and if a user is added to channel j, all users currently on channel j are turned off if they can no longer achieve C_{target}.

We note that when all active users achieve the same rate Jain fairness is equal to the fraction of active users. Similarly, sum-rate per user is equal to the target-rate multiplied by the fraction of active users. Therefore, because both plots show similar information, we will neglect to show Jain fairness. Instead, when we compare results for target based systems, we show a histogram of the number of active users. In Figure 3.23, we show the histogram of the number of active users for $N = 6$ and $M = 2$. We can see that when we have a lower target rate, it is more likely that more users are active (i.e., for a target of 1 bps all users are active about 5% of the time); however, when we have a higher target rate, it is more likely that more users are cut off (i.e., for a target of 5 bps, no users are active about 10% of the time). In Figure 3.24, we plot the fraction of active users against the target rate. We notice again that by increasing the target rate, fewer users are active. However, from Figure 3.25 we can see that as we increase the target rate the average sum-rate per user increases. We expect the average sum-rate per user to increase as the target rate increases from zero because higher capacities are possible; however, as we increase the target rate past some threshold, the average sum-rate per user will begin to decrease because fewer and fewer users will be able to utilize the available channels (due to fewer users being able to achieve the target rate). As the target rate increases, the fairness steadily decreases due to more users not achieving the target rate.

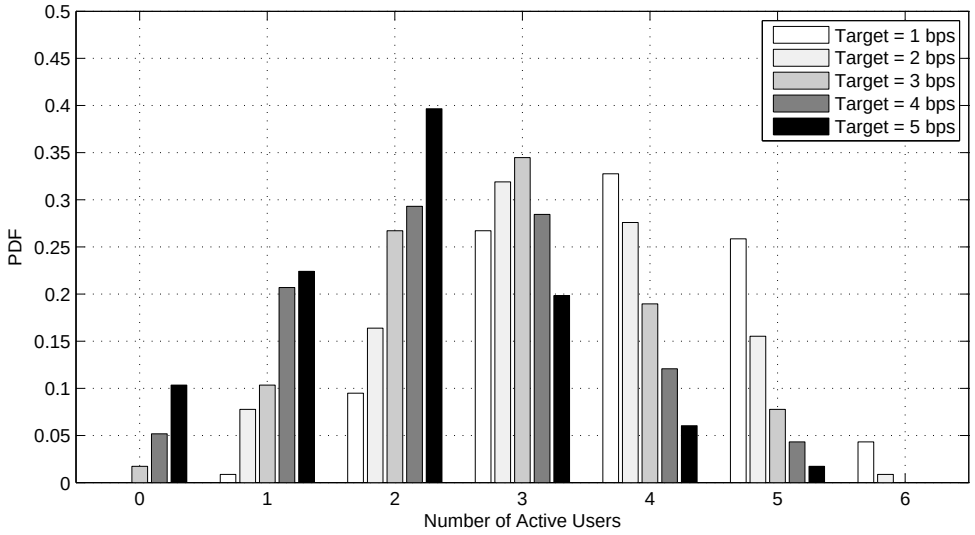

Figure 3.23: Histogram of Number of Active Users with Varying Target Rates ($SNR_{min} = 3dB$, $MUI = 1$, $N = 6$ and $M = 2$)

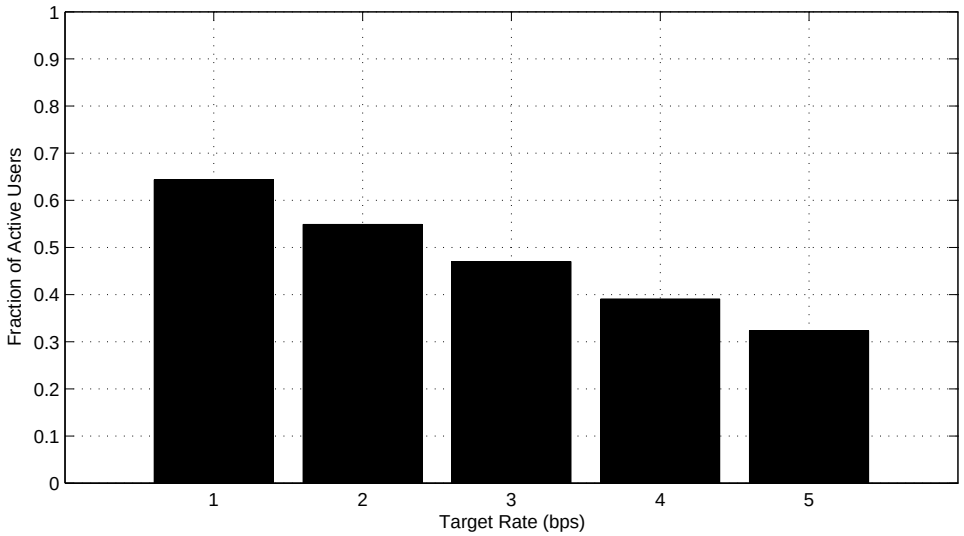

Figure 3.24: Fraction of Active Users for Varying Target Rates ($SNR_{min} = 3dB$, $MUI = 1$, $N = 6$ and $M = 2$)

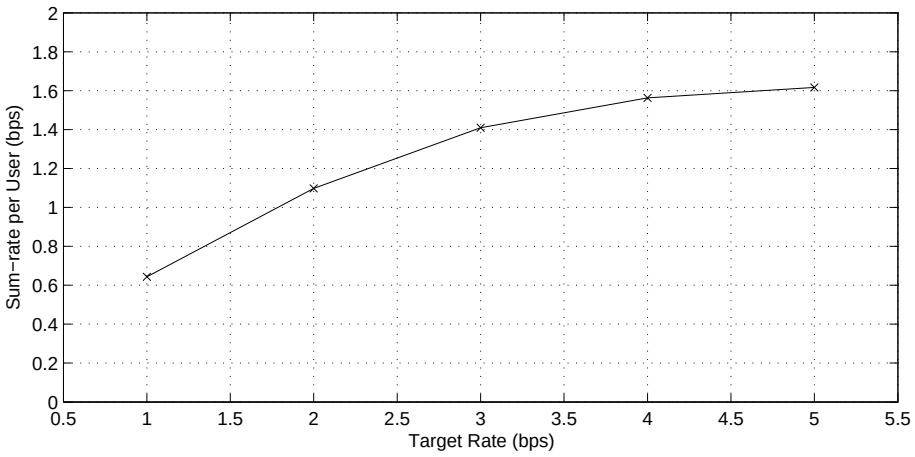

Figure 3.25: Average Sum-rate per User for Varying Target Rates ($SNR_{min} = 3dB$, $MUI = 1$, $N = 6$ and $M = 2$)

3.7 Conclusions

Based upon the work in this chapter, we make the following conclusions:

- Adding users to a system causes the sum-rate per user and Jain fairness to decrease. We note that while it decreases as is expected, it does not follow the trend of Gupta and Kumar in [1] because sensing is not optimal.
- Adding channels to a system causes the sum-rate per user and Jain fairness to increase.
- Using multi-channel sensing tends to yield a sum-rate per user benefit, but this benefit comes at the cost of losing fairness by further enhancing the performance of the users with the best capacities while hurting the performance of the users with the worst capacities.
- The benefit of using multi-channel sensing over single-channel sensing increases when the number of users decreases and when the number of channels increase.
- Using multiple rounds of sensing tends to benefit both the sum-rate per user and the Jain fairness simultaneously of both SC and MC sensing by increasing the capacities of the worst users.
- Increasing interference power by placing users onto the plane more densely quickly decreases performance while increasing noise power steadily decreases performance.
- When using a target-based sensing approach, increasing the target rate results in decreased fairness.
- Increasing the target rate past zero causes the sum-rate per user to increase, but the sum-rate per user reaches a maximum and begins to decrease when the target rate increases past the maximum due to more users being cutoff.

In the next chapter, we will introduce our model for full collaboration and show some initial simulation results to compare the performance of full collaboration to no collaboration.

Chapter 4

Modeling Collaboration

In the previous chapter, we showed the performance of non-collaborative approaches. In this chapter, we consider the performance of collaborative approaches and begin to compare the collaborative and non-collaborative results. In Section 4.1, we introduce the two optimization problems (corresponding to the two different objective functions that we consider in this work). We use multiple objective functions because collaboration can be used to improve capacity, fairness, or increase the number of users that can achieve a desired target rate. In Section 4.2, we introduce the local search algorithm that we use to solve the optimization problems; since the results of a local search algorithm are not guaranteed to be globally optimal, we verify the proximity of these results to the global optimum. Finally, we show simulation results for optimizing sum-rate, optimizing with fairness considerations, and optimizing the number of admitted users with a target-based approach.

4.1 Optimization Problem

In this section, we introduce the optimization problems that we use to model full collaboration. We begin by presenting our two main objective functions followed by a few system constraints (those constraints inherent to practical modulation schemes, i.e. maximum bits per second (bps) limit). We also introduce some optimization constraints; these are the constraints that we add to the optimization routine to attempt to have a certain outcome (for example if we want to achieve at least a particular fairness level while still maximizing another objective function).

4.1.1 Sum-rate Objective Function

To optimize the sum-rate of the system, we maximize the sum of the individual user rates (as shown in (2.16)). Optimization is performed simultaneously over channel selection and transmit power. The optimization problem used to maximize sum-rate is given in (4.1). Constraints of the channel access techniques will be given in Section 4.1.3.

$$\begin{array}{ll} \max & \sum_{i=1}^{N} C_i \\ \text{s.t.} & \sum_{k=1}^{M} P_{i,k} \leq P_{max} \\ & \text{Constraints of the Channel Access Technique} \end{array} \tag{4.1}$$

4.1.2 Sum of Log-rate Objective Function

A second objective used in this work (besides sum-rate) is sum of log-rate. This is often referred to as *proportional fairness* in the literature. (The equivalence of sum of log-rate and the traditional definition of proportional fairness is explored in [91]. It is shown that maximizing sum of log-rate yields the proportionally fair solution.) Proportional fairness provides a solution that balances capacity and fairness by ensuring that no user has zero capacity while allowing the users with better channels to have higher capacities. More specifically, proportional fairness is achieved by maximizing the objective function related to sum of log-rate:

$$\sum_{i=1}^{N} log\left(C_i\right) \tag{4.2}$$

Note that proportional fairness avoids the situation of any user having a capacity of zero because this makes the sum of log-rate go to negative infinity. The optimization problem that maximizes sum of log-rate is given in (4.3). Constraints of the channel access techniques will be given in Section 4.1.3.

$$\begin{array}{ll} \max & \sum_{i=1}^{N} log_2\left(C_i\right) \\ \text{s.t.} & \sum_{k=1}^{M} P_{i,k} \leq P_{max} \\ & \text{Constraints of the Channel Access Technique} \end{array} \tag{4.3}$$

4.1.3 System Constraints

This section introduces the system constraints that we may enforce that are inherent to real world scenarios. We will begin by explaining the constraints necessary for each channel access technique. We will then introduce limitations on capacity due to coding and modulation techniques and limitations on achievable $SINR$ to account for our inability to obtain Shannon capacity in real world scenarios.

Channel Access Technique Constraints

We begin by noting that all channel access techniques are limited by a maximum transmit power per user of P_{max}. MC and SS have no additional constraints beyond the power constraint. However, for the case of the SC access technique, each user is limited to transmitting on a maximum of one channel. For all simulations, we assume SC unless otherwise stated. We note that using SC makes the optimization problem a non-linear mixed-integer optimization problem (it is mixed-integer because channel selection is discrete while power is continuous).

Modulation and Coding Constraints

All channel access techniques are also limited by the modulation and coding limitations given in Chapter 2. These include a maximum spectral efficiency, a minimum spectral efficiency, and an $SINR$ multiplier.

4.1.4 Optimization Constraints

Here we focus on those constraints that we add during the optimization to attempt to achieve some goal (i.e., better fairness) while optimizing a separate objective function. These constraints are implemented through Lagrange multiplier techniques. While our constraints make the optimization attempt to achieve the goal, it is possible that the constraint is infeasible and the optimization fails.

Jain Goal

The first optimization constraint we consider is a Jain fairness constraint where Jain fairness is given by (2.18). We use this to attain different levels of fairness while optimizing the sum-rate of the system. The value multiplied by the Lagrange multiplier is given in (4.4).

$$Jain_{goal} - Jain_{actual} \tag{4.4}$$

Minimum Rate Goal

The second optimization constraint we consider is a constraint on the minimum user capacity. The value multiplied by the Lagrange multiplier is given in (4.5).

$$C_{goal} - C_i \tag{4.5}$$

This constraint is similar to the global constraint on the minimum capacity, but it differs in two ways. First, the global constraint is implemented for each user on each channel, while the optimization constraint is implemented on the capacity across all channels for each user (this only makes a difference in MC). Second, the optimization constraint on minimum capacity is implemented through a Lagrange multiplier whereas the global constraint on minimum capacity is implemented through a hard limit.

4.2 Approaches for finding / bounding the 'optimal'

To solve the non-linear mixed-integer optimization problem associated with using SC techniques, we use the branch and bound (BB) technique as is commonly used in similar problems [30, 92]. However, the BB technique will not always return the global optimum; therefore, we must determine how reliable the solution of the BB technique is. The result of the BB optimization depends upon the initial solution given to it; therefore, by running the optimization multiple times with different initial solutions and selecting the best result, we can improve the results of the BB technique. For example, if the BB fails to return the global optimum for 10% of the initial solutions that we generate and we perform the optimization twice with two different arbitrary initial solutions, we only have a 1% probability of failing to find the global optimum with the combined result.

To determine the reliability of the branch and bound technique, we must compare the solution of the BB technique with the true global optimum across the entire search space. The technique we use to find the true global optimum is the branch and bound relaxation-linearization technique (BB-RLT) [93]. We note that, while the BB-RLT algorithm is guaranteed to find results within a small percentage of the true global optimum, the computational complexity of this algorithm is orders of magnitude larger than the computational complexity of the BB technique (for some scenarios where the BB technique takes around 15 seconds, the BB-RLT technique can take many hours). Because the time consuming nature of the BB-RLT makes it implausible for running thousands of optimizations, we would like to use the BB technique to perform our optimizations.

We begin this section by showing the methodology behind using the BB-RLT to find the global optimum. We then use the BB-RLT technique to find the global optimum for a set of scenarios and simulate the same scenarios with BB to determine the closeness of the BB solutions to the global optimum. Specifically, we determine how many optimizations with different initial solutions are necessary to ensure that the performance of the BB is within 1% of the global optimum found by the BB-RLT.

4.2.1 Finding the Global Optimum with BB-RLT

The BB-RLT method finds the global optimum by creating a new optimization problem that relaxes the constraints of the original non-linear problem in such a way that the new problem is linear. Since the linear optimization problem only relaxes the non-linear terms and encompasses the feasible region of the original problem, the solution is an upper bound (UB) on the optimum of the non-linear optimization problem. We define the lower bound (LB) to be the best solution of the non-linear optimization problem that has been found (this value may increase as the optimization solves the non-linear problem with different bounding intervals in each iteration). The algorithm works in such a way as to continuously partition the search space into smaller search spaces such that the LB and the largest UB both approach the global optimum. When the largest UB and the LB are within some threshold, ϵ, the BB-RLT is complete and yields a LB on the global optimum that is within ϵ of the true global optimum. For this work, we use $\epsilon = 0.1\%$.
The BB-RLT process is summarized in Table 4.1 (this table is taken from [2] with slight modifications). In the following discussion, we describe how we use the BB-RLT for our optimization problem.

Restating the Non-linear Optimization Problem

The optimization problem we are attempting to solve is defined as shown in 2.12. Because we are trying to reformulate this into a linear optimization problem, we will rewrite it in a way that looks more linear as shown in (4.6). Assuming we use SC, the variables being optimized over include the continuous power variables, P_i, and binary channel selection variables, $S_{i,j}$. The $S_{i,j}$ values take values of either 0 or 1 and the P_i variables range from $0 \leq P_i \leq P_{max}$ (for this analysis we assume $P_{max} = 1$). The power transmitted by user i on channel j is then equal to $P_i \cdot S_{i,j}$.

$$\begin{array}{ll} max & \sum_{i=1}^{N} \sum_{j=1}^{M} \frac{1}{ln(2)} \cdot (Y_{i,j} - W_{i,j}) \\ s.t. & Y_{i,j} = ln\,(X_{i,j}) \\ & W_{i,j} = ln\,(V_{i,j}) \\ & X_{i,j} = \sum_{k=1}^{N} g_{k,i,j} \cdot T_{k,j} + \nu \\ & V_{i,j} = X_{i,j} - g_{i,i,j} \cdot T_{i,j} \\ & T_{i,j} = P_i \cdot S_{i,j} \\ & S_{i,j} = 0, 1 \\ & \sum_{j=1}^{M} S_{i,j} \leq 1 \\ & 0 \leq P_i \leq 1 \end{array} \tag{4.6}$$

We note that by writing the problem like this, we have added one new variable for each non-linear term as is required in Step 4 of the Initialization stage, namely we have added $Y_{i,j}$, $W_{i,j}$, and $T_{i,j}$.

Table 4.1: BB-RLT Technique [2]

Algorithm 1 BB/RLT Solution Procedure

Initialization:

1. Let optimal solution $\psi^* = \theta$ and the initial lower bound $LB = -\infty$.
2. Determine partitioning variables and derive their initial bounding intervals.
3. Let the initial problem list contain only the original problem, denoted by P_1.
4. Introduce one new variable for each nonlinear term. Add linear constraints for these variables to build a linear relaxation. Denote the solution to linear relaxation as $\hat{\psi}_1$ and its objective value as the upper bound UB_1.

Main Loop:

1. Select problem P_z that has the largest upper bound among all problems in the problem list.
2. Find a feasible solution, ψ_z, via a local search algorithm for Problem P_z. Denote the objective value of ψ_z by LB_z.
3. If $LB_z > LB$ then let $\psi^* = \psi_z$ and $LB = LB_z$. If $LB \geq (1-\epsilon) \cdot UB_z$ then stop with the $(1-\epsilon)$-optimal solution ψ^*; else, remove all problems $P_{z'}$ having $(1-\epsilon) \cdot UB_{z'} \leq LB$ from the problem list.
4. Intelligently select a partitioning variable and divide its bounding interval into two new intervals.
5. Remove the selected problem P_z from the problem list, construct two new problems P_{z1} and P_{z2} based on the two partitioned intervals.
6. Compute two new upper bounds UB_{z1} and UB_{z2} by solving the linear relaxations of P_{z1} and P_{z2}, respectively.
7. If $LB < (1-\epsilon) \cdot UB_{z1}$ then add problem P_{z1} to the problem list. If $LB < (1-\epsilon) \cdot UB_{z2}$ then add problem P_{z2} to the problem list.
8. If the problem list is empty, stop with the $(1-\epsilon)$-optimal solution ψ^*. Otherwise, repeat the Main Loop.

Add Linear Constraints for the New Non-linear Variables

Because Step 4 of the Initialization requires that we add linear constraints for the non-linear variables, we add constraints for $T_{i,j}$, $Y_{i,j}$ and $W_{i,j}$. The non-linearity of $Y_{i,j}$ and $W_{i,j}$ are due to the presence of the natural log function while the non-linearity of $T_{i,j}$ is due to the multiplication of two optimization variables. We will bound each of these respective to the variables involved in the non-linearity.
Because $T_{i,j} = P_i \cdot S_{i,j}$ and we know both P_i and $S_{i,j}$ are non-negative, we can bound $T_{i,j}$ as shown in (4.7). These bounds are recalculated for each subproblem.

$$\begin{array}{rcl} T_{i,j,UB} & = & P_{i,UB} \cdot S_{i,j,UB} \\ T_{i,j,LB} & = & P_{i,LB} \cdot S_{i,j,LB} \end{array} \tag{4.7}$$

To bound Y and W, we use the technique shown in [2]. Figure 4.1 shows the polyhedral approximation to the natural log function that is used to bound Y based upon the value of X. This outer approximation is made by determining X_{LB} and X_{UB} based upon the values of the T_{LB} and T_{UB} and making four linear constraints about the relationship between $Y_{i,j}$ and $X_{i,j}$. To form a tighter bound, a new variable $X_{LB} < X_B < X_{UB}$ is defined as shown in (4.8). The four bounds are then defined as shown in (4.9). Similar linear constraints are made in the relationship between W and V. These bounds are recalculated for each subproblem.

$$X_B = \frac{X_{UB} \cdot X_{LB} \cdot (ln\,(X_{UB}) - ln\,(X_{LB}))}{X_{UB} - X_{LB}} \tag{4.8}$$

$$\begin{array}{lrcl} I. & X_{LB} \cdot y - x & \leq & X_{LB} \cdot (ln\,(X_{LB}) - 1) \\ II. & X_B \cdot y - x & \leq & X_B \cdot (ln\,(X_B) - 1) \\ III. & X_{UB} \cdot y - x & \leq & X_{UB} \cdot (ln\,(X_{UB}) - 1) \\ IV. & (X_{UB} - X_{LB}) \cdot y + (ln\,(X_{LB}) - ln\,(X_{UB})) \cdot x & \geq & X_{UB} \cdot ln\,(X_{LB}) - X_{LB} \cdot ln\,(X_{UB}) \end{array} \tag{4.9}$$

State the Linear Relaxation of the Original Problem

With all of the non-linear terms relaxed and linearly constrained, we now formulate the linear optimization problem.

$$\begin{array}{ll} max & \sum_{i=1}^{N} \sum_{j=1}^{M} \frac{1}{ln(2)} \cdot (Y_{i,j} - W_{i,j}) \\ s.t. & \text{Four linear constraints on } (Y_{i,j}, X_{i,j}) \\ & \text{Four linear constraints on } (W_{i,j}, V_{i,j}) \\ & X_{i,j} = \sum_{k=1}^{N} g_{k,i,j} \cdot T_{k,j} + \nu \\ & V_{i,j} = X_{i,j} - g_{i,i,j} \cdot T_{i,j} \\ & \text{Bounding constraints for } T_{i,j} \end{array} \tag{4.10}$$

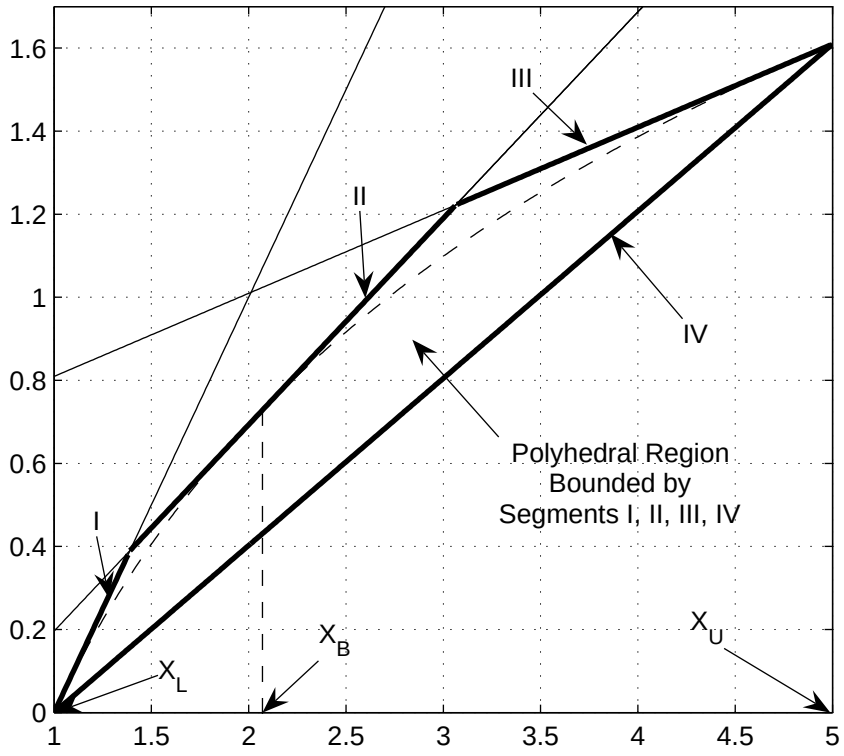

Figure 4.1: Polyhedral Outer-Approximation for $y = ln\,(x)$

Partitioning

The one remaining step in the main loop that we have not described in detail is the selection of a partitioning variable and the partitioning of that variable. To partition the continuous P_i variables, we split the range $P_{LB} \leq P \leq P_{UB}$ into $P_{LB} \leq P \leq \frac{P_{LB}+P_{UB}}{2}$ and $\frac{P_{LB}+P_{UB}}{2} \leq P \leq P_{UB}$. To partition the binary $S_{i,j}$ variable, we split the range $0 \leq S \leq 1$ into $0 \leq S \leq 0$ and $1 \leq S \leq 1$. By partitioning the binary variables this way, we enforce the requirement that the $S_{i,j}$ variables be binary without creating a non-linearity in the relaxed optimization problem.

Selection of the partitioning variable and partitioning technique is left up to the implementor but should be done intelligently to minimize the time required to run the optimization algorithm. We found that for our application partitioning on binary variables first is most efficient. Selection of which binary variable to partition on is arbitrary. Once all binary variables for a subproblem have been partitioned, the power variables are partitioned. Selection

of the power variable to partition is arbitrary amongst the power variables with the largest remaining range (i.e., those power variables with the largest value of $P_{UB} - P_{LB}$).

4.2.2 Reliability of Branch and Bound

To determine the reliability of the original simpler branch and bound procedure, we create individual scenarios in which we determine the global optimum using BB-RLT and the potentially sub-optimal result of the BB technique and compare the two. For all simulations, we set $SNR_{min} = 5dB$ and $MUI = 5$. We measured the sub-optimality for the parameter cases of $N = 1$ and $M = 3$, $N = 4$ and $M = 1$ and $N = 5$ and $M = 3$. For each of these parameter cases, we created 200 specific scenarios and determined the global optimum using BB-RLT. For each scenario, we run the BB optimization with a fixed number of arbitrary initial solutions.

To measure optimality, we plot a CDF curve of the sum-rate of the BB-RLT and a CDF curve of the sum-rate using the BB technique (assuming 1 initial solution unless otherwise stated). We determine the percentage of sub-optimality at 20%, 50%, and 80% on the CDF curve. Our goal is to determine the number of BB optimizations with different initial solutions that are necessary to ensure that on average the result of the BB optimization is within 1% of the global optimum.

In Figure 4.2, we see that for the case of four users and a single channel the global optimum is around 7.1 and the BB results in around 6.9 at the 50% line (percentage decrease of around 3%). For the case of a single user and three channels, there is almost no loss due to using branch and bound ($< 1\%$) as seen in Figure 4.3. In Figure 4.4, we can see in the CDF of sum-rate for the case of $N = 5$ and $M = 3$ that the median sum-rate of the global optimum is around 20.4 while the median sum-rate resulting from the branch and bound technique is around 19.8, a loss of about 3%. We note that adding channels seems to have very little impact on the optimality of the branch and bound, but as we add more users, the ratio of the branch and bound result to the global optimum decreases.

To enable our error rate to average around 1%, we perform BB optimization multiple times using different initial solutions and selecting the result that is largest. In Figures 4.5 and 4.6, we repeat the simulation using $N = 5$ and $M = 3$ using two and three initial solutions respectively. By adding more initial solutions it can be seen that the suboptimal BB technique and the global optimum become closer. In Figure 4.7, we repeat the $N = 4$ and $M = 1$ parameters using three initial solutions. In this situation, we see the same results. The measure of suboptimality of the BB technique for all scenarios we considered is summarized in Table 4.2.2. Since three initial solutions seems to give results within 1% of the global

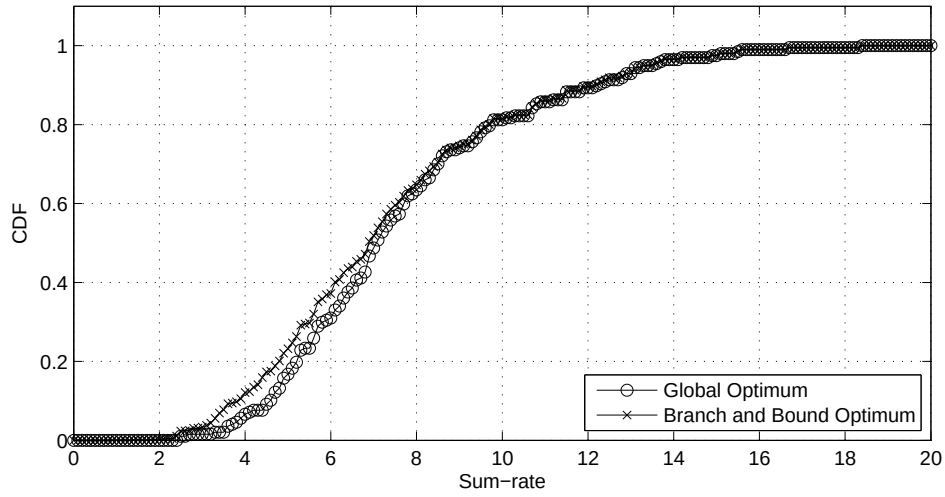

Figure 4.2: Comparison of the Global Optimum vs. the Optimum Found by the Branch and Bound Procedure (M=4, N=1)

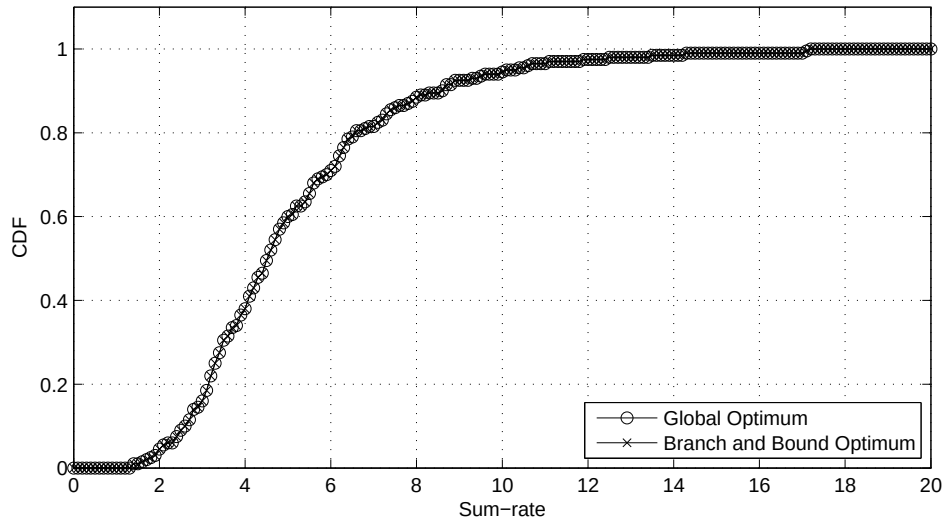

Figure 4.3: Comparison of the Global Optimum vs. the Optimum Found by the Branch and Bound Procedure (M=1, N=3)

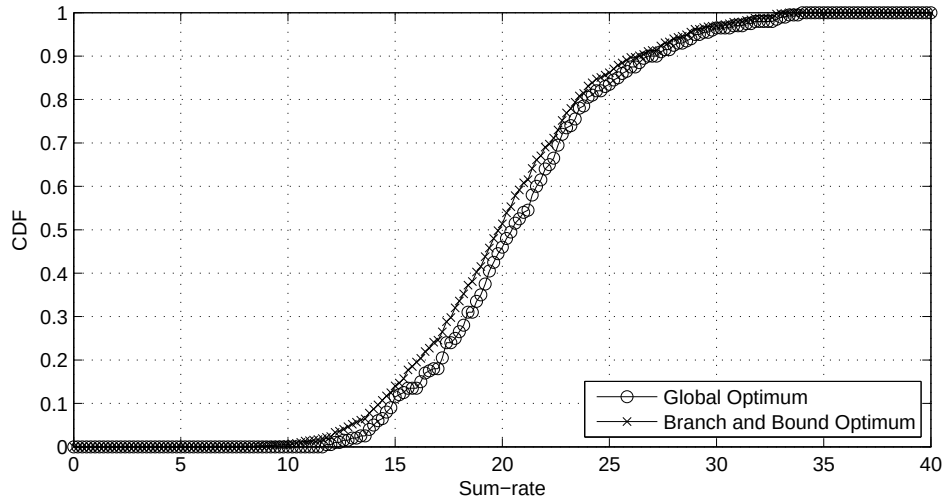

Figure 4.4: Comparison of the Global Optimum vs. the Optimum Found by the Branch and Bound Procedure (M=5, N=3)

optimal for our test cases, we will use three initial solutions for our simulation routines when using BB throughout this work.

Table 4.2: Percentage Drop from Globally Optimal Sum-rate from Using Branch and Bound

CDF Percentage Line	20%	50%	80%
N=4 M=1 (1 Initial Solution)	7.5%	3%	<1%
N=4 M=1 (3 Initial Solutions)	2%	1%	<1%
N=1 M=3 (1 Initial Solution)	<1%	<1%	<1%
N=5 M=3 (1 Initial Solution)	6%	3%	2%
N=5 M=3 (2 Initial Solutions)	3%	1.5%	1%
N=5 M=3 (3 Initial Solutions)	1%	1%	<1%

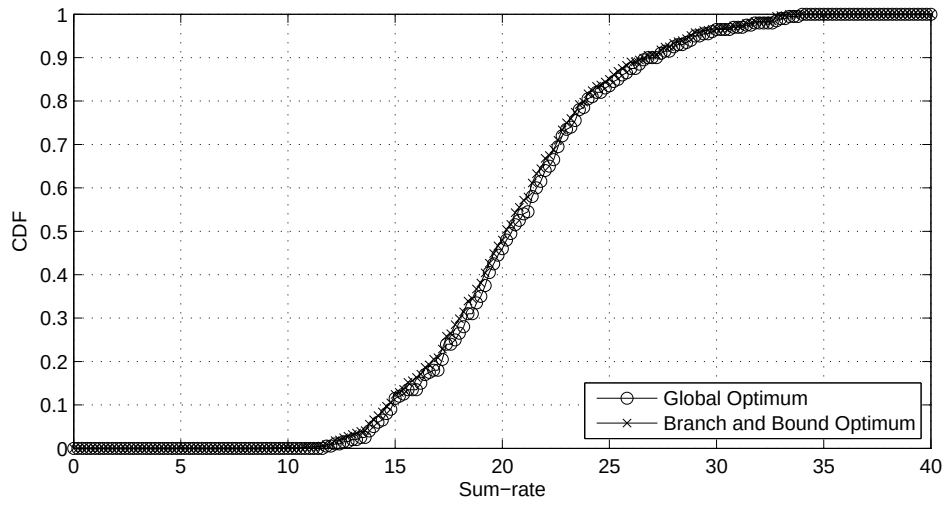

Figure 4.5: Comparison of the Global Optimum vs. the Optimum Found by the Branch and Bound Procedure (M=5, N=3, Initial Solutions=2)

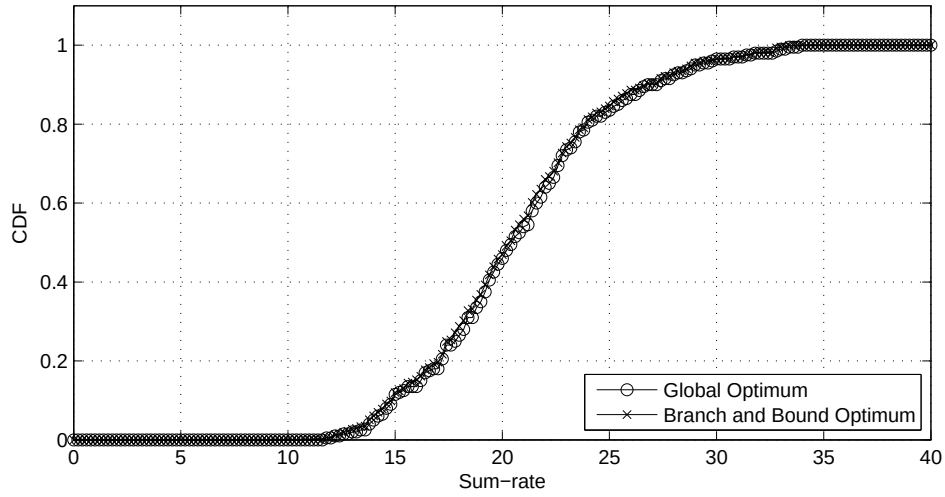

Figure 4.6: Comparison of the Global Optimum vs. the Optimum Found by the Branch and Bound Procedure (M=5, N=3, Initial Solutions=3)

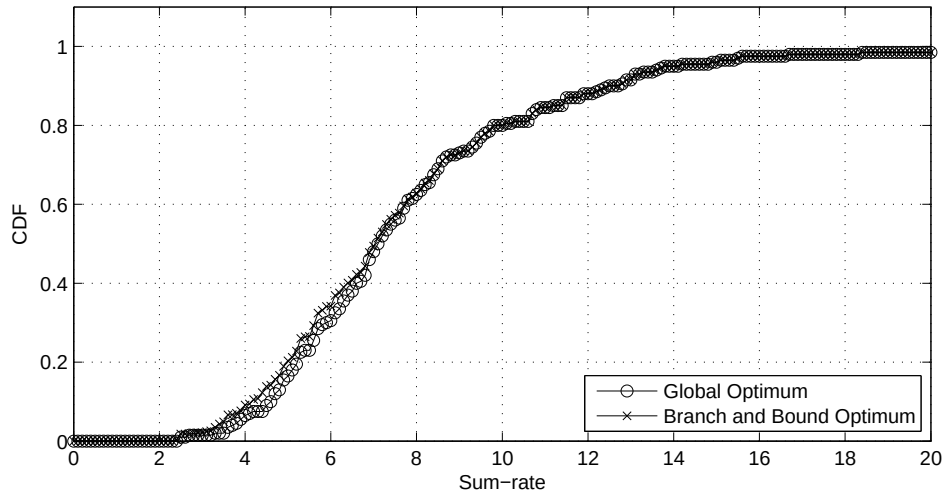

Figure 4.7: Comparison of the Global Optimum vs. the Optimum Found by the Branch and Bound Procedure (M=4, N=1, Initial Solutions=3)

4.3 Optimizing Capacity

In this section, we compare simulation results to determine the impact of different system parameters on both the achievable sum-rate per user for SC optimization and the benefit of collaboration. These include the amount of shadowing in the environment, the predictability of the fading, the level of the noise power, the closeness of interfering users, the number of users, the number of channels, whether outside interference is present, whether links are uni-directional or bi-directional, and having multiple optimization groups (aka. partial collaboration).

4.3.1 Sample CDF Result and the Collaboration Gain

In this section, we present an example case to illustrate how we compare the results using CDF plots and how we determine the collaboration gain. We simulate the case in which 7 users and 2 channels are present with no outside interference. We assume an SNR_{min} of $3dB$ and a MUI of 5 and plot the CDF of the sum-rate per user in Figure 4.8. From the figure, we can see that SC optimization yields a performance improvement of about 50% over sensing (median performance increases from about 1.8 bps to about 2.8 bps). The average sum-rate per user for sensing is around 2 bps and the average sum-rate per user for SC optimization is around 3 bps. For future reference, we will define the term *collaboration gain* to be the gain of the average sum-rate per user of SC over the average sum-rate per user of sensing (in our case this is $\frac{3}{2} = 1.5$).

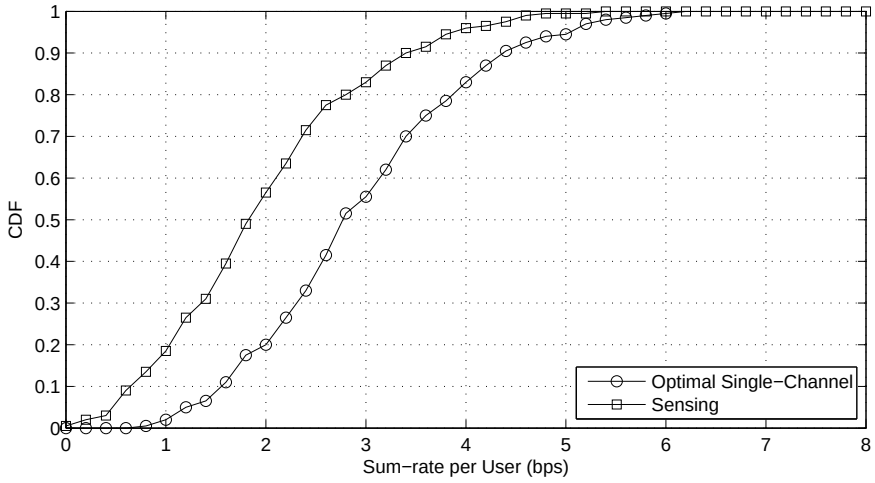

Figure 4.8: Sample CDF of Sum-rate per User for 7 Users 2 Channels ($SNR_{min} = 3dB$ and $MUI = 5$)

4.3.2 Impact of Channel Gain

In Section 2.3, we introduced the channel gain assumptions that we use for our simulations. We assume that the channel gain term is the product of a path loss term, a shadowing term, and a multipath fading term. In this section, we first compare the simulation results for different levels of shadowing. After this, we compare the simulation results for slow and fast fading. We make conclusions about the impact of each on the benefit of collaboration and state the assumed values upon which we perform all remaining simulations.

Using an SNR_{min} of $3dB$ and a MUI of 5, we plot the relationship between average sum-rate per user of SC optimization and sensing for different shadowing parameters in Figure 4.9. All values are plotted against the number of users and parameterized by the number of channels. The collaboration gain is plotted in the first row, the average sum-rate per user for SC optimization is plotted in the second row, and the average sum-rate per user for sensing is plotted in the third row. The plots in the left column relate to lower shadowing effects (a shadowing coefficient of 1), and the plots in the right column relate to higher shadowing effects (a shadowing coefficient of 3). Comparing (a) to (b), we can see that as the shadowing increases, the benefit of collaboration also increases. The increase in the shadowing coefficient causes the average sum-rate per user of both optimization ((c) and (d)) and sensing ((e) and (f)) to increase. These gains can be attributed to two factors. First, as we increase the shadowing, the expected channel gains increase because the shadowing is log-normal distributed. This effect evenly impacts both SC optimization and sensing because the impact effectively decreases the expected noise power. (We note that in real scenarios an increase in shadowing should reduce channel gains; however, since will maintain a constant shadowing factor throughout the remainder of this report and since we are mainly concerned with the impact of shadowing on collaboration gains, this should not harm our results since it impacts SC optimization and sensing evenly.) However, SC optimization achieves a further benefit of multi-user diversity. The multi-user diversity results from the increase in the variance in the channel gains due to the increased variance of the shadowing. This makes it more likely that the user with the best channel in a high shadowing environment will have a better channel gain than the user with the best channel in a low shadowing environment (assuming all parameters except shadowing remain constant). This has a larger impact on the optimization because, as we will show later, optimization tends to favor transmission for the users with the best channel gains and disadvantage users with the worst channel gains. We use a shadowing factor of 3 for all of the following simulations. This means that the results we show on the benefit of collaboration are pessimistic if actual shadowing effects are higher and optimistic if the actual shadowing effects are lower.

Recall that because our optimization is performed once for each system setup (and thus does not consider time), the speed of the fading does not affect the optimization routine itself. Instead, the speed of the fading impacts the channel gains that we optimize over. We assume

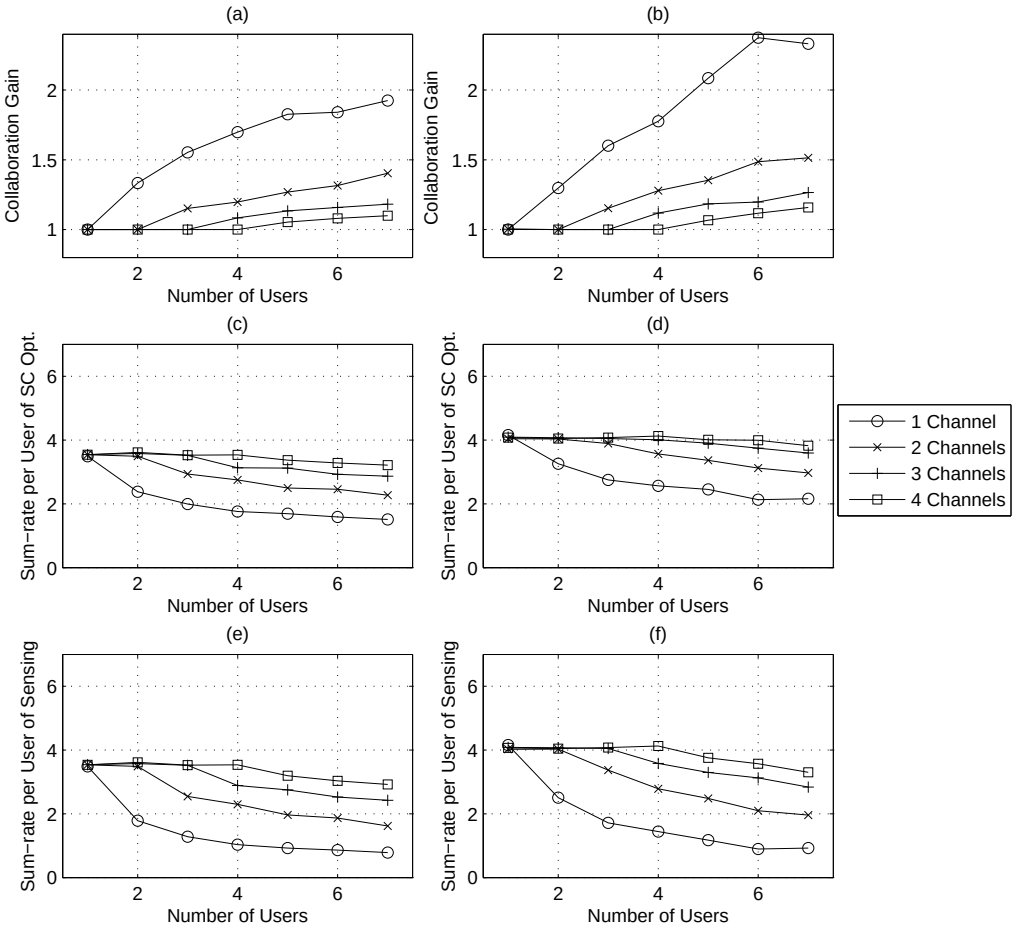

Figure 4.9: Impact of Shadowing Parameter. The left column has less shadowing ($\sigma = 1$) and the right column has more shadowing ($\sigma = 3$). The top row shows the collaboration gain, the middle row shows the sum-rate per user of SC optimization, and the bottom row shows the sum-rate per user of sensing.

that if fast fading is present, the fading changes so quickly that it is unpredictable; in this case, optimization is performed over channel gain terms consisting of only path loss and shadowing (estimating the channel by averaging over time removes the fading term). But if slow fading is present, the fading is slow enough that we assume it is predictable, and the optimization is performed over the channel gain terms consisting of path loss, shadowing, and fading. For both slow and fast fading, actual system performance is determined based upon the channel gain terms consisting of path loss, shadowing, and fading.

Using an SNR_{min} of $3dB$ and a MUI of 5, we plot the relationship between average sum-rate per user for SC optimization and sensing of both slow and fast fading in Figure 4.10. All values are plotted against the number of users and channels. The collaboration gain is plotted in the first row, the average sum-rate per user for SC optimization is plotted in the second row, and the average sum-rate per user for sensing is plotted in the third row. The plots in the left column relate to slow fading, and the plots in the right column relate to fast fading. We first note in Figure 4.10 that the benefit of collaboration is largely unchanged from (a) to (b). From looking at the average sum-rate per user plots in (c)-(f), we can see slow, predictable fading yields a significant benefit to the average sum-rate per user of both SC and sensing when multiple channels are present. We attribute this benefit to the multi-channel diversity that results from fading being predictable (this makes the channel gains different across the channels). The sum-rate per user curves nearly match for the case in which a single channel is present because multi-channel diversity cannot yield any benefit when only a single-channel exists. We can see in (d) and (f) that, when fading is fast and $N \leq M$, SC optimization and sensing have identical performance since the channel gains are perfectly correlated across all channels (preventing channel selection from benefiting any user). In these cases, optimization and sensing perform identically since they both randomly place users on different channels. Therefore, since the multi-channel diversity affects both SC optimization and sensing almost evenly, the collaboration gains of both techniques are nearly identical. Because of this, we assume slow fading for all remaining simulations. If fading is actually fast, the collaboration gains we find should still be accurate, but the average sum-rate per user values will be optimistic.

4.3.3 Impact of Users and Channels

In this section, we explore the impact of the number of users and channels on the average sum-rate per user of SC optimization. For the case of $SNR_{min} = 3dB$ and $MUI = 5$, we have plotted the average sum-rate per user for SC optimization in Figure 4.11. We notice that as the number of channels increases the average sum-rate per user also increases; however, this benefit follows the law of diminishing returns as adding a third channel gives a lesser benefit than adding a second channel. The benefits are caused by a combination of an increase in the amount of bandwidth available (more channels) and an increase in multi-channel diversity. To illustrate why this benefit follows the law of diminishing returns consider the following

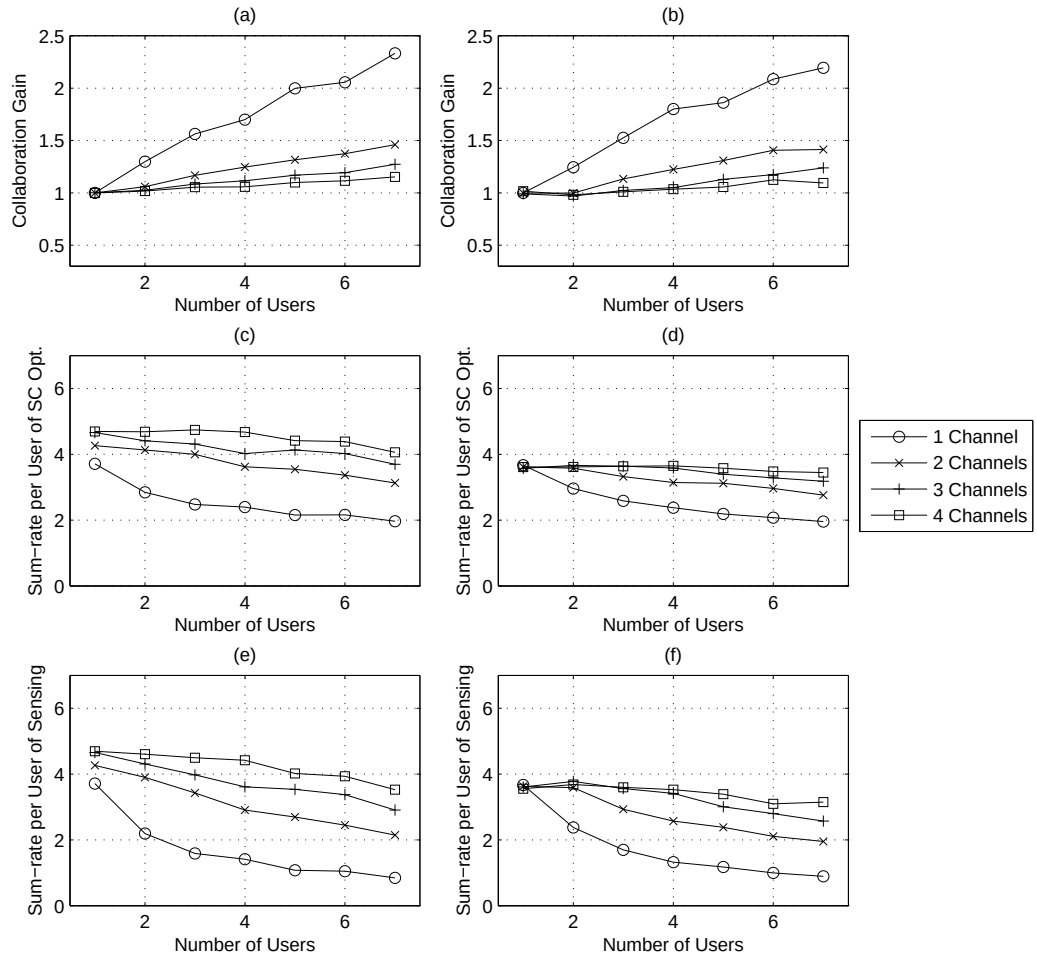

Figure 4.10: Impact of Slow vs. Fast Fading. The left column is slow fading and the right column is fast fading. The top row shows the collaboration gain, the middle row shows the sum-rate per user of SC optimization, and the bottom row shows the sum-rate per user of sensing. ($\sigma = 3$)

example. If the set of users that achieve the optimal sum-rate with a channel are given priority to that channel and a second channel is added, it is unlikely that a subset of the remaining users would be able to achieve as high of a capacity on the second channel that was achieved on the original channel. In this case, the capacity increase resulting from the second channel is generally less than the capacity on the original channel because additional channels tend to be given to sets of users with progressively worse performance. From the figure, we also notice that as the number of users increases, the average sum-rate per user decreases. Two factors affect this decline, namely a small increase due to multi-user diversity and a decrease due to the sum-rate per user being divided by a larger number of users. The multi-user diversity results from an increase in the expected channel gain of the best users (by order statistics) and more freedom for the optimization routine to select which users transmit as the number of users increases. Both of these result in an increased efficiency in channel usage as the number of users increases. However, this benefit is overpowered by the decrease due to dividing the sum-rate by an additional user (if we plot the sum-rate instead of the sum-rate per user, the curve consistently increases due to the multi-user diversity benefit). Considering this another way, the expected capacity of a new user should be less than the expected capacity of current users in a system since this new user increases the amount of interference present in the system.

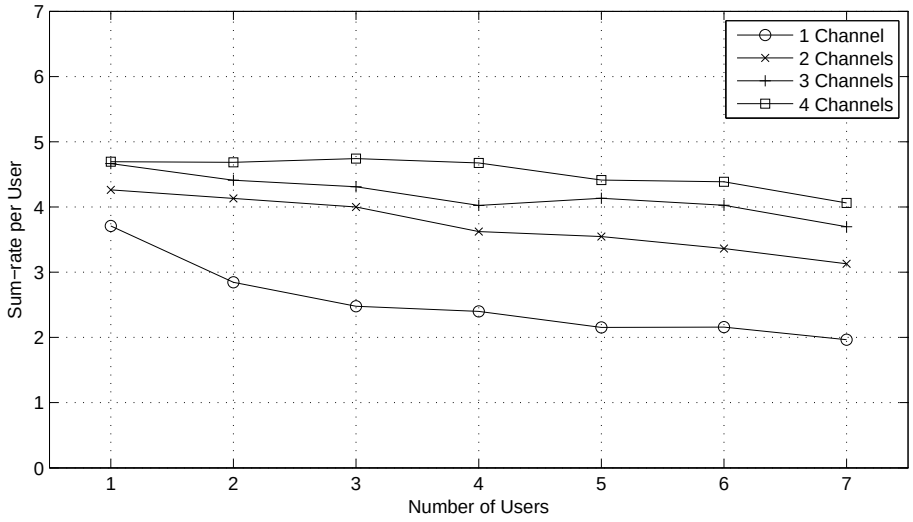

Figure 4.11: Impact of Number of Users and Channels on Average Sum-rate per User ($SNR_{min} = 3dB$ and $MUI = 5$)

In Figure 4.12, we can see the value of the collaboration gain for different numbers of users and channels. Because the benefit of collaboration is interference avoidance, the collaboration gain increases as we increase the number of users and thus increase the interference. Similarly, as we increase the number of channels, the expected interference power decreases causing the benefit of collaboration to decrease.

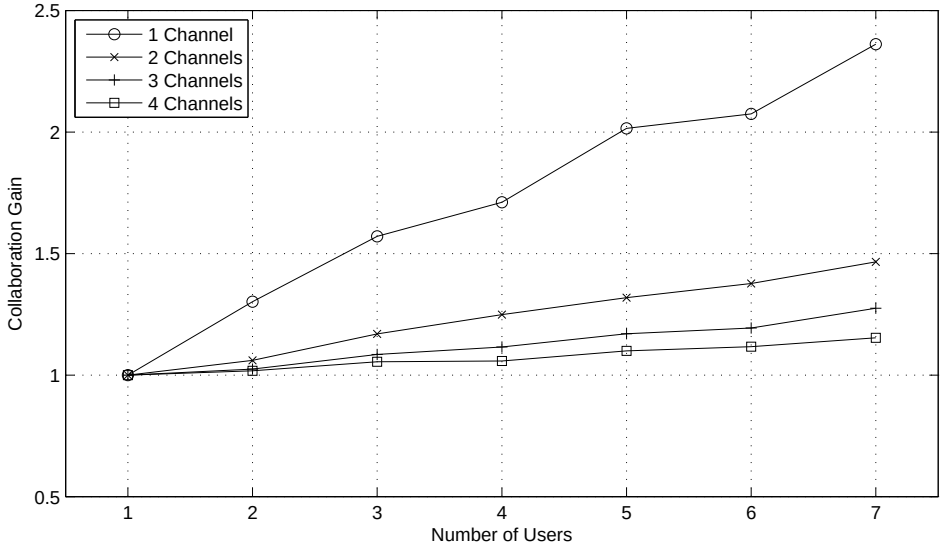

Figure 4.12: Impact of Number of Users and Channels on Collaboration Gain ($SNR_{min} = 3dB$ and $MUI = 5$)

We note that a theoretical trend of the capacity per user as a function of the number of users has been derived by Gupta and Kumar in [1] for a system model similar to our own. The main differences between our system model and that of Gupta and Kumar is that as the number of users increases, our user density remains constant (for Gupta and Kumar it increases), and we place users inside a square while Gupta and Kumar place users inside a circle. From the definition of MUI, we note that our model should nearly match that of Gupta and Kumar when the number of users is equal to the MUI (by nearly, we mean that the only difference remaining is that we use a square while Gupta and Kumar use a circle). When $N < MUI$, the interference for our model should be higher than the interference for the model of Gupta and Kumar, so our average sum-rate per user for this scenario should be lower than the theoretical value. When $N > MUI$, the interference for our model should be lower than the interference for the model of Gupta and Kumar, so our average sum-rate per user for this scenario should be higher than the theoretical value.

The results from Gupta and Kumar say that the expected capacity per user for N arbitrarily placed users decays with $\frac{1}{N^{\frac{1}{\eta}}}$ (where η is the path loss factor which is 3 for us). We plot the simulated average sum-rate per user with the theoretical expected capacity curve in Figure 4.13 (for the theoretical capacity curve, we plot $\frac{C_{1-User}}{N^{\frac{1}{3}}}$ where C_{1-User} is the capacity with only a single user present). We see that our simulation results are very near to the theoretical value. As expected, when $N < MUI$ the performance of our simulations are lower than the theoretical value and about when $N = MUI = 5$ our simulation results are nearly equal to

the theoretical value; however the difference is very minor.

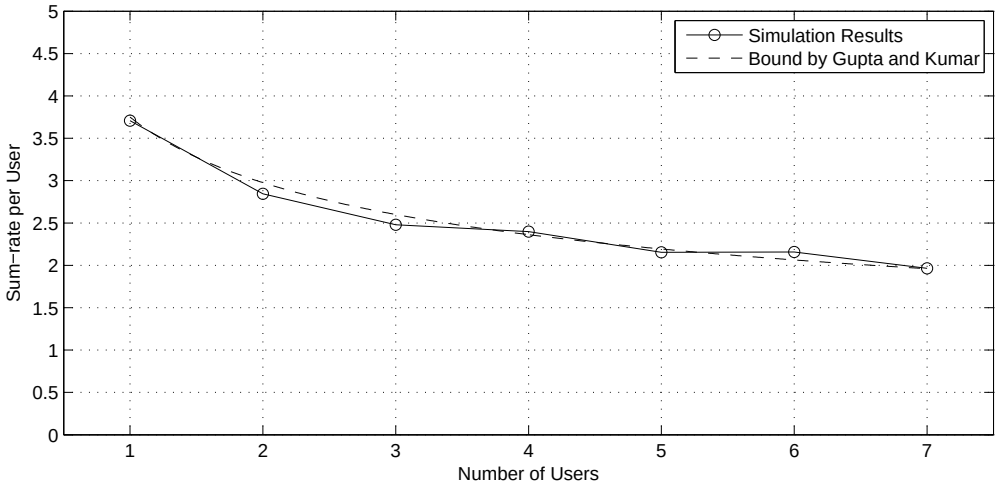

Figure 4.13: Comparison of Simulation Results to Theoretical Value from [1] ($MUI = 5$, $SNR_{min} = 3dB$, $M = 1$)

We now compare the theoretical bound of Gupta and Kumar to the case in which we have two channels. Because the bound is intended to be for users on a single channel, we modify our axis to be number of users divided by number of channels. Therefore, our theoretical bound will decay with $\left(\frac{M}{N}\right)^{\frac{1}{3}}$. We plot the simulation results along with the bound in Figure 4.14 (we adjust the bound to the capacity of two users). We can see that simulation results for the two channel case decays at a much slower rate than the theoretical bound. This is due to the multi-channel diversity that allows users to separate onto two channels if they would strongly interfere with each other on a single channel. The second channel causes users to be more likely to be able to avoid interference by avoiding channels where harmful interference may be caused.

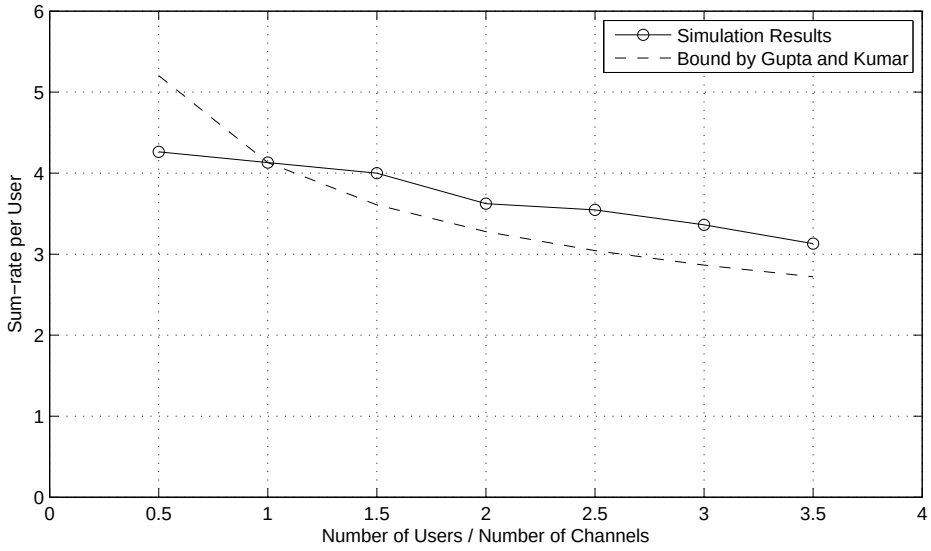

Figure 4.14: Comparison of Simulation Results to Theoretical Value from [1] ($MUI = 5$, $SNR_{min} = 3dB$, $M = 2$)

4.3.4 Impact of Noise and Interference

In Figure 4.15, we explore the impact of interference and noise on the benefit of collaboration. The y-axis on the plot is the collaboration gain (as explained in 4.3.1). We can see that as the MUI (or interference) increases (or expected SIR decreases), the benefit of collaboration increases. Similarly, as the SNR_{min} increases (or noise power decreases), the benefit of collaboration increases. These two parameters are related in that when noise is the dominating factor in the $SINR$ denominator, the benefit of collaboration is smaller, whereas when interference is the dominating factor in the $SINR$ denominator, the benefit of collaboration is larger. This is what we would expect because collaboration is intended to minimize interference, so if we are noise-limited and interference does not significantly affect capacity, the benefit of collaboration should be minimal. However, if we are interference-limited, collaboration should help us to minimize that interference, so the benefit of collaboration should be higher.

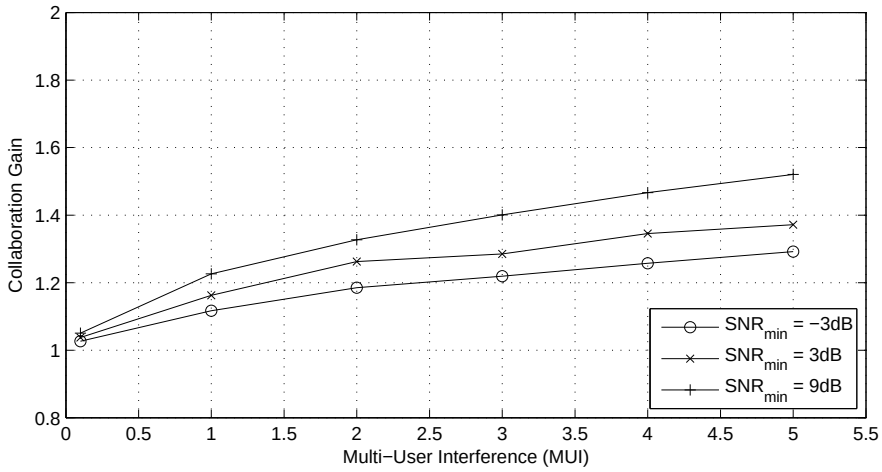

Figure 4.15: Impact of Interference and Noise on the Collaboration Gain ($N = 6$ and $M = 2$)

In Figure 4.16, we explore the impact of noise and interference on the average sum-rate per user of full collaboration. We can see that as the SNR_{min} increases the average sum-rate per user increases for a fixed value of MUI. Increasing the SNR_{min} by 6dB seems to cause the average sum-rate per user to increase by around 1 bps for most values of MUI shown on the plot. As the MUI increases, the sum-rate per user decreases as expected (increased MUI means users are more likely to cause harmful interference to each other). The largest decrease in average sum-rate per user occurs when the MUI increases from 0.1 to 1. This is likely because at a MUI of 0.1 most users can transmit simultaneously without causing each other harmful interference; however when MUI equals 1, it is likely that users will begin to interfere with each other. As MUI increases, we expect the average sum-rate per user to

approach a constant which is equal to the sum of the capacities of the best two users (each user receives its own channel). As we will show later, the users with the best channels will be likely to have the majority of the capacity, so the average sum-rate per user does not decrease significantly as the MUI increases from 0.1 to 5.

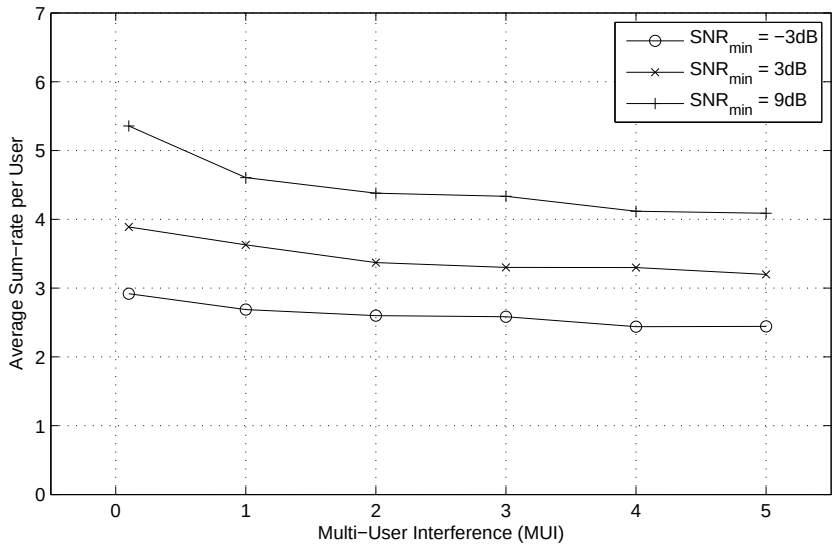

Figure 4.16: Impact of Interference and Noise on the Average Sum-rate per User of Full Collaboration ($N = 6$ and $M = 2$)

4.3.5 Impact of Outside Interference

In this section, we determine the impact of the outside interference loading factor on the collaboration gain (the outside interference loading factor was introduced in Section 2.2). The collaboration gain is plotted against the outside interference factor for a MUI of 1 and 5 in Figure 4.17 for the case of $N = 6$ and $M = 2$. We can see that adding an outside interference loading of 0.1 causes the collaboration gain to decrease from 1.37 to 1.22 for $MUI = 5$ and from 1.16 to 1.08 for $MUI = 1$. This is because addition of outside interference essentially results in an increased noise floor since this interference cannot be controlled. As the outside interference loading factor increases, we quickly become limited by this outside interference which has the same effect as being limited by noise (the benefit of collaboration quickly decreases).

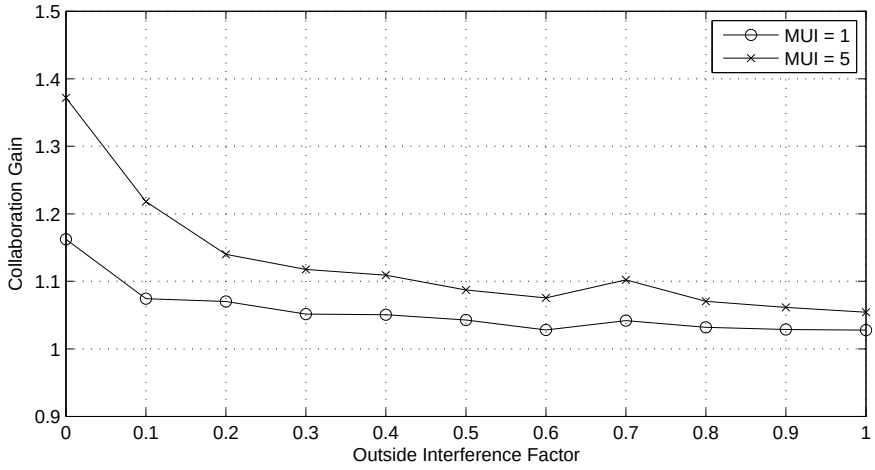

Figure 4.17: Impact of Outside Interference ($SNR_{min} = 3dB$, $N = 6$ and $M = 2$)

4.3.6 Impact of Reversing Transmission

Because communications are not typically uni-directional, we consider the impact of bi-directional links on the results of the optimization routine. Specifically, we optimize the channel selection and power values assuming a uni-directional link and determine whether reversing the transmissions of a given fraction of the users (which we will call the *reversal fraction*) while maintaining the optimal power and channel values negatively impacts the sum-rate per user or the collaboration gain. We plotted the collaboration gain against the expected fraction of users who reverse transmission (reversal fraction) in Figure 4.18, the average sum-rate per user of optimizing sum-rate against the reversal fraction in Figure 4.19, and the average sum-rate per user of sensing against the reversal fraction in Figure 4.20. As

the reversal fraction approaches 0.5, the sum-rate per user of optimizing sum-rate tends to decrease but only slightly. As the reversal fraction approaches 1, the sum-rate per user tends to increase back to previously higher levels. A similar effect can be seen in the collaboration gain and the sum-rate of sensing (although for the sum-rate of sensing this effect is much smaller. The maximum decrease in the collaboration gain (where the reversal fraction is 0.5) is quite small (around 0.05) for both $MUI = 1$ and $MUI = 5$. Increasing the MUI to 5 causes the impact of the reversal fraction to be more severe on the sum-rate per user values of both optimal sum-rate and sensing. We plot the average sum-rate per user of optimizing proportional fairness in Figure 4.21. We can see that, while reversing transmission does not have as large of an impact on sum-rate per user as it does when optimizing sum-rate, it still has an impact similar to these other examples (best performance at a reversal fraction of 0 or 1 and worst at 0.5).

In both optimization and sensing, Tx nodes tend to transmit on channels that give them high gains to their corresponding Rx node, but small gains from other Tx nodes to their own Rx node (interference gains). We assume in this analysis that the gain from Tx node a to Rx node b is the same as the gain from Rx node b to Tx node a. Since the channel gain from a Tx node to the corresponding Rx node is constant (regardless of reversal), reversing transmission only causes the interference gains to change. When complete reversal occurs (reversal fraction = 1), the gain from Tx node a to Rx node b in the original problem is now the gain from Tx node b to Rx node a. Since the original channel selection and power control were aimed to minimize the interference gains, the interference gains are still minimized in the complete reversal because the same set of interference gains are used. However, when we use a reversal fraction of 0.5, new interference gains are included in the capacity calculation that were not considered in the original optimization or sensing procedure (i.e., Tx node a to Tx node b or Rx node a to Rx node b was not optimized in the original problem).

We consider an example to show how this may impact the sum-rate per user by the scenario shown in Figure 4.22 where we assume that users 1 and 2 share the same channel. For simplicity, we assume that channel gains are calculated by $g_{i,j} = \frac{1}{d_{i,j}^2}$. This means that $g_{1,1} = g_{2,2} = 1$ and $g_{1,2} = g_{2,1} = 0.4$; therefore, if we assume maximum transmit power of 1 and no noise, the SIR for each will be 2.5, and the individual capacities will be 1.8 bps (thus sum-rate is 3.6 bps). In the case of complete reversal, the gains would be exactly the same, so achievable sum-rate would be the same (we note that if $g_{1,2} \neq g_{2,1}$ this would not be the case, however the interference gains would still be minimal in the problem with complete reversal). If user 2 reverses transmission, $g_{1,1} = g_{2,2} = 1$, $g_{1,2} = 0.22$, and $g_{2,1} = 2$. This means $C_1 = 0.58bps$ and $C_2 = 2.17bps$, so the sum-rate would be $C = 2.75bps$. This drop in sum-rate results from the unexpectedly high interference gain from Rx node 1 to Rx node 2 that is not considered when placing users 1 and 2 on the same channel in the original sensing procedure or optimization problem. When reversing all users or no users, the interference gains are the same set of interference gains that are considered in the original problem (and

this is why reversing all users and no users both achieve nearly optimal performance). When reversing half of the users, we introduce the most of the interference gains that were not considered in the original problem (which is why a reversal fraction of 0.5 tends to yield the worst performance.

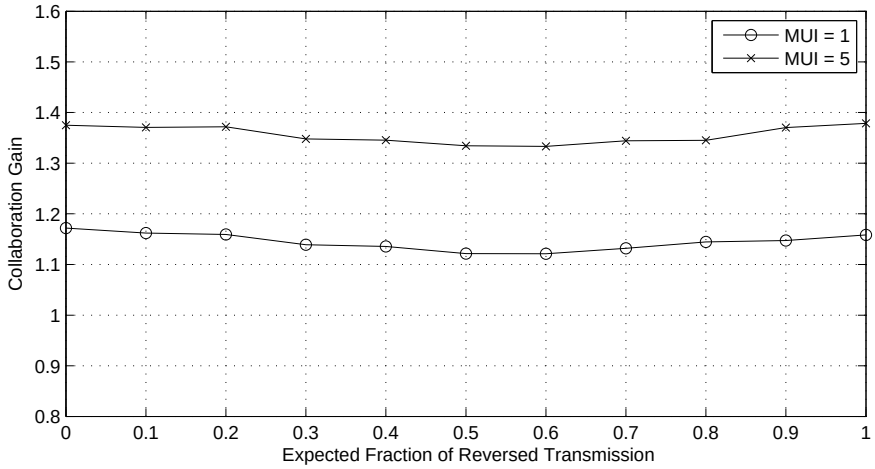

Figure 4.18: Impact of Reversing Transmission on Collaboration Gain ($SNR_{min} = 3dB$, $N = 6$ and $M = 2$)

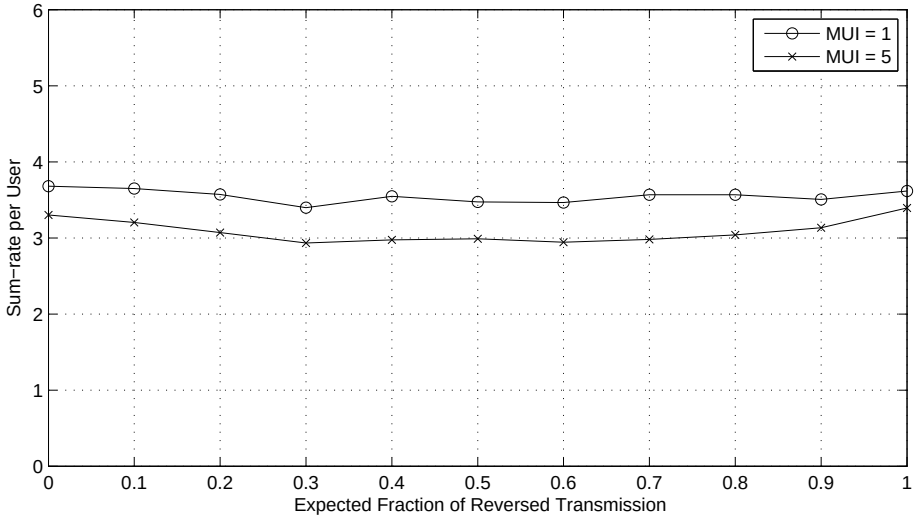

Figure 4.19: Impact of Reversing Transmission on Sum-rate per User of Optimizing Sum-rate ($SNR_{min} = 3dB$, $N = 6$ and $M = 2$)

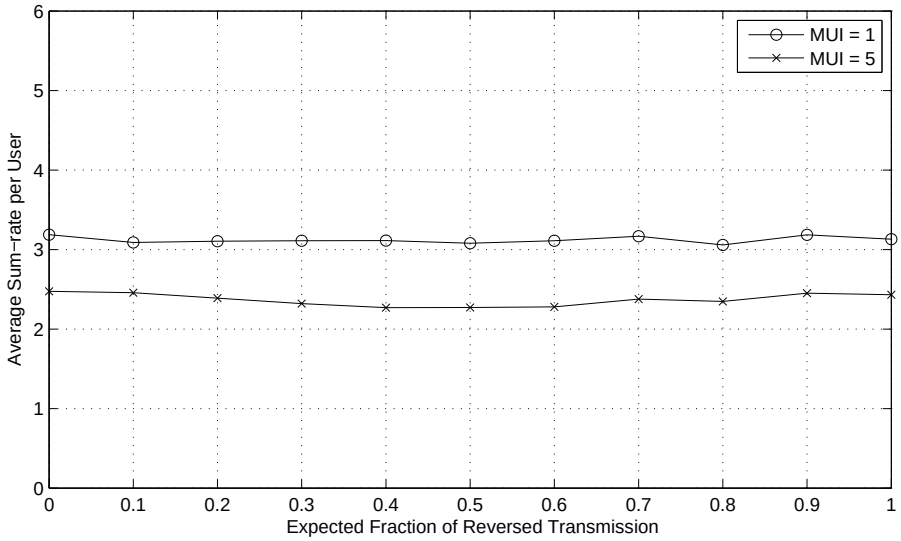

Figure 4.20: Impact of Reversing Transmission on Sum-rate per User of Sensing ($SNR_{min} = 3dB$, $N = 6$ and $M = 2$)

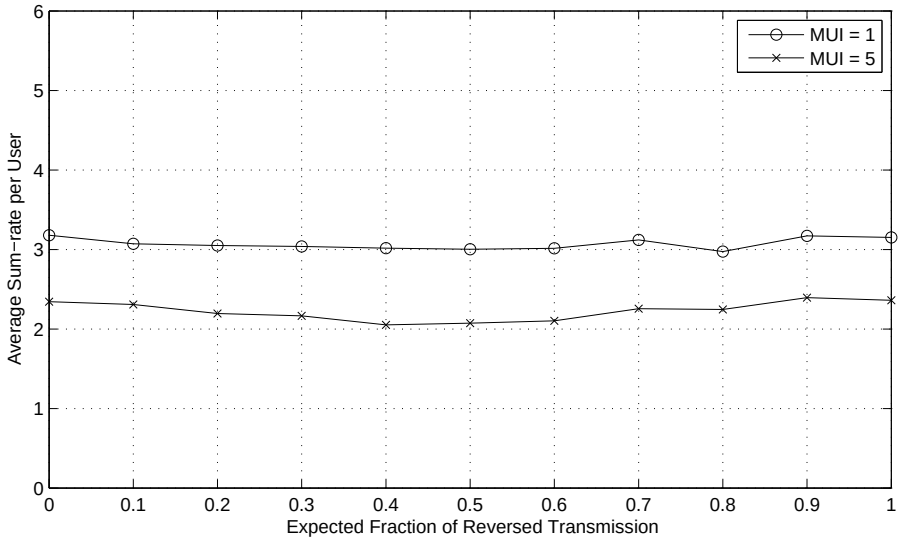

Figure 4.21: Impact of Reversing Transmission on Sum-rate per User of Optimizing Sum of Log-rate ($SNR_{min} = 3dB$, $N = 6$ and $M = 2$)

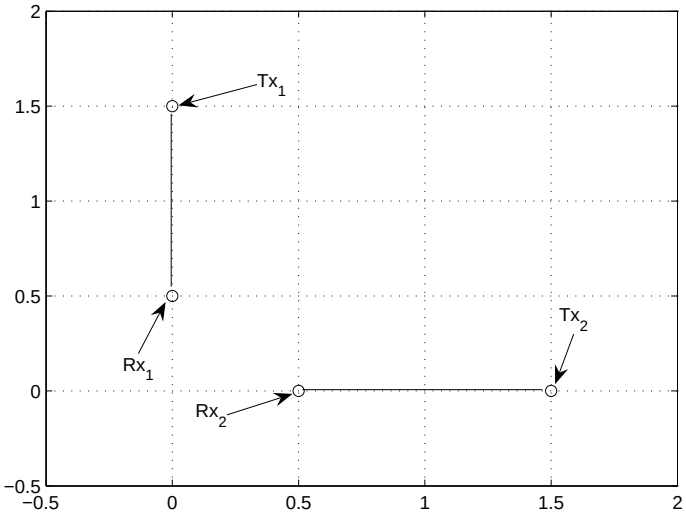

Figure 4.22: Sample Case of the Impact of Reversing Transmission ($SNR_{min} = 3dB$, $N = 6$ and $M = 2$)

4.3.7 Impact of Partial Collaboration

In this section, we study the impact of partial collaboration. Specifically, we focus on the scenario mentioned in Section 2.5 where users may be willing but unable to fully collaborate. Users are placed into groups by using the k-means algorithm described in Section 2.5 and groups are added sequentially with each group only optimizing amongst the users within the group. For the scenario where $MUI = 5$, $SNR_{min} = 3dB$, we ran simulations for 3 to 7 users with the number of groups ranging from 1 to the number of users. We plot the collaboration gain against the number of groups for 2 channels in Figure 4.23 and 3 channels in Figure 4.24. We note that one group is equivalent to full collaboration and when the number of groups and users are the same this is equivalent to sensing.

As expected, we can see in both plots that the benefit of collaboration decreases as the number of groups increases. Interestingly, the slope of the curves tend to be nearly linearly decreasing as the number of groups increases; however, the slopes tend to flatten out as the collaboration gain approaches 1.

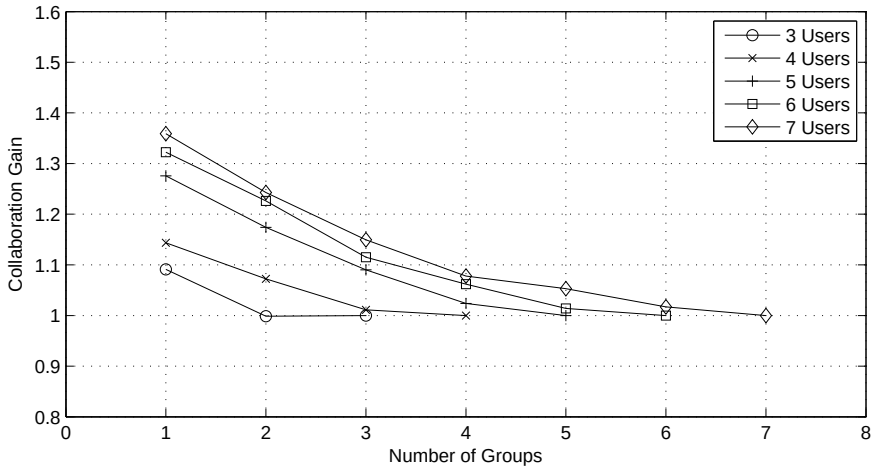

Figure 4.23: Impact of Limited Collaboration ($MUI = 5$, $SNR_{min} = 3dB$, and $M = 2$)

4.3.8 Impact of Modulation and Coding Limitations

We now consider the impact of the modulation and coding limitations introduced in Chapter 2 on the sum-rate per user and collaboration gain. We begin by noting that our optimization routine requires that the optimization parameters be continuous because it uses a gradient technique to solve the optimization problem. Because the minimum bps/Hz limitation causes a discontinuity (for example when a user's power passes a threshold the capacity of that user

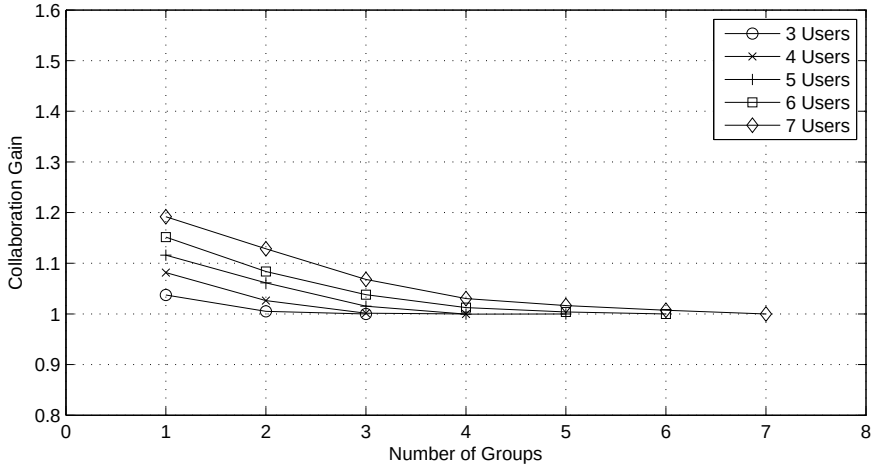

Figure 4.24: Impact of Reversing Transmission ($MUI = 5$, $SNR_{min} = 3dB$, and $M = 3$)

jumps from zero to the minimum capacity), we cannot guarantee the proximity of results in which a minimum bps/Hz limitation is implemented to the global optimum. By using the minimum bps/Hz optimization constraint, we can maximize the number of users that achieve the minimum bps/Hz, but this does not maximize sum-rate. Therefore, we will not study the minimum bps/Hz limitation when optimizing sum-rate because it is not evident how we would achieve such a result. We will consider in Chapter 6 how a minimum bps/Hz limitation affects the performance of sensing.

For the maximum bps/Hz limitation, we plot the performance of both full collaboration and sensing for limits of 6 bps/Hz, 8 bps/Hz and no limitation in Figure 4.25. We can see that the performance degrades as we implement the maximum spectral efficiency limits; this is expected because we essentially are removing freedom as we decrease the maximum spectral efficiency. We note that the collaboration gain for the 6 bps/Hz limit is 1.376, the collaboration gain for the 8 bps/Hz limit is 1.378, and the collaboration gain for no limitation is 1.374. This means that the maximum bps/Hz limitation for reasonable upper limits does not impact the collaboration gain.

We plot the performance of both full collaboration and sensing for $SINR$ multipliers (ξ) of 0.5, 0.9 and 1 in Figure 4.26. We can see that as the $SINR$ multipliers decrease the performance of both techniques degrade nearly equivalently (by about 0.5 bps for $\xi = 0.5$). Users with the highest values of $SINR$ are expected to have a degradation of around 1 bps for $\xi = 0.5$; users with the lowest $SINR$ values are expected to have very low losses (nearly zero for some users) for $\xi = 0.5$. Combined with our performance loss of around 0.5 bps, these mean that the best users to suffer significant performance loss, but the worst

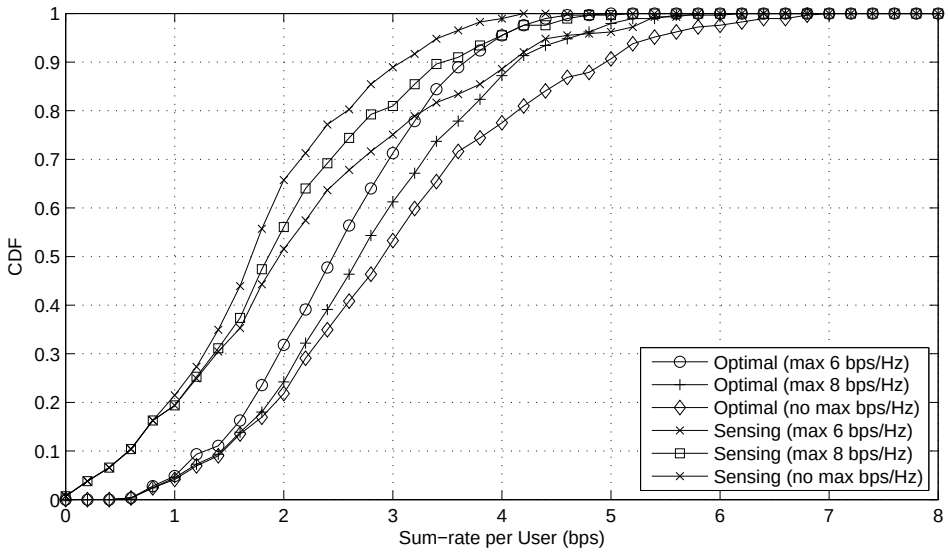

Figure 4.25: Impact of a Maximum Spectral Efficiency on Sum-rate per User ($MUI = 5$, $SNR_{min} = 3dB$, $N = 6$, $M = 2$)

users tend to suffer very little loss. We note that the collaboration gain for $\xi = 1$ is 1.374, the collaboration gain for $\xi = 0.9$ is 1.377, and the collaboration gain for $\xi = 0.5$ is 1.424. Thus as the $SINR$ multiplier decreases the collaboration gain increases, but this increase in the collaboration gain is very small for a large degradation in performance. Therefore, the $SINR$ multiplier has very little impact on the collaboration gain.

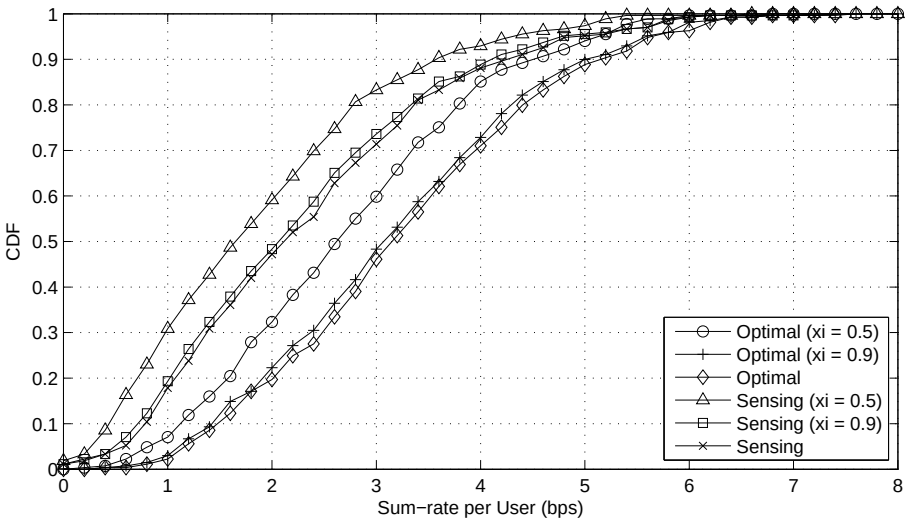

Figure 4.26: Impact of $SINR$ Multiplier ($MUI = 5$, $SNR_{min} = 3dB$, $N = 6$, $M = 2$)

4.4 Optimizing for Fairness

For the case of $N = 6$ and $M = 2$, we plot the CDF of the individual capacities comparing optimal sum-rate and sensing for the best two users in Figure 4.27, the middle two users in Figure 4.28, and the worst two users in Figure 4.29. We notice that the capacities of the users with the best capacities and the users with the middle capacities both increase when we optimize sum-rate. However, the capacities of the users with the worst capacities decrease when we optimize sum-rate. Therefore, when we optimize sum-rate the users with the best capacities are given preference over the users with the worst capacities. This leads to uneven fairness of the spectral resources amongst users. Because fairness is an important consideration in wireless systems, we consider the impact of collaboration on fairness in this chapter.

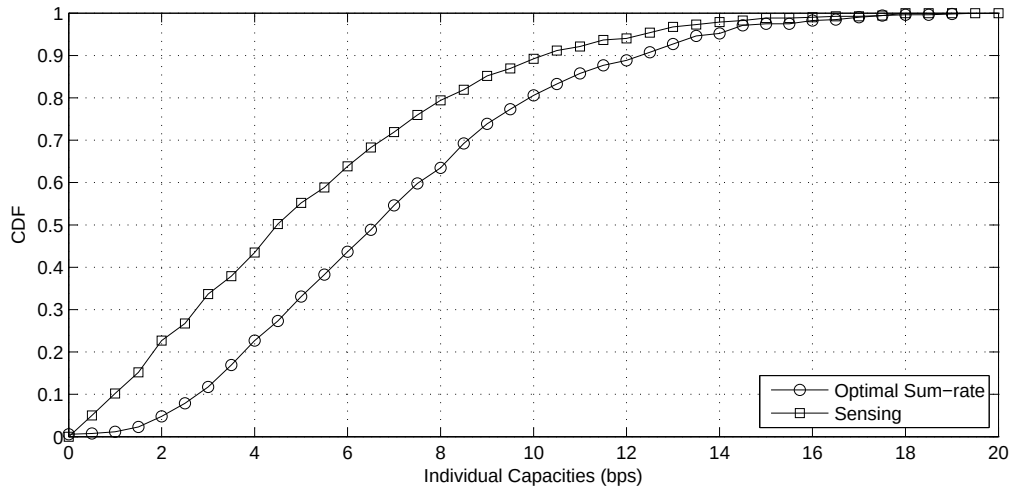

Figure 4.27: CDF of Individual Capacities of the Best Two Users ($MUI = 1$, $SNR_{min} = 3dB$, $N = 6$, and $M = 2$)

In this section, we consider a few alternative optimization problems that balance fairness and capacity (unlike optimizing sum-rate which increases capacity with no regard to fairness). Specifically, we compare optimizing sum-rate with optimizing sum of log-rate and optimizing sum-rate with a Jain fairness optimization constraint goal of 0.5. In the comparisons, we include sensing because, while sensing might not achieve optimal sum-rate, it may still achieve good fairness.

To compare the four techniques, we plot the CDF of the sum-rate per user for the case where $SNR_{min} = 3dB$, $N = 6$, and $M = 2$ using an MUI of 1 in Figure 4.30 and a MUI of 5 in Figure 4.31. We first note that optimizing sum-rate yields the highest sum-rate per user (as it obviously should). Optimizing sum of log-rate and sensing tend to share the worst

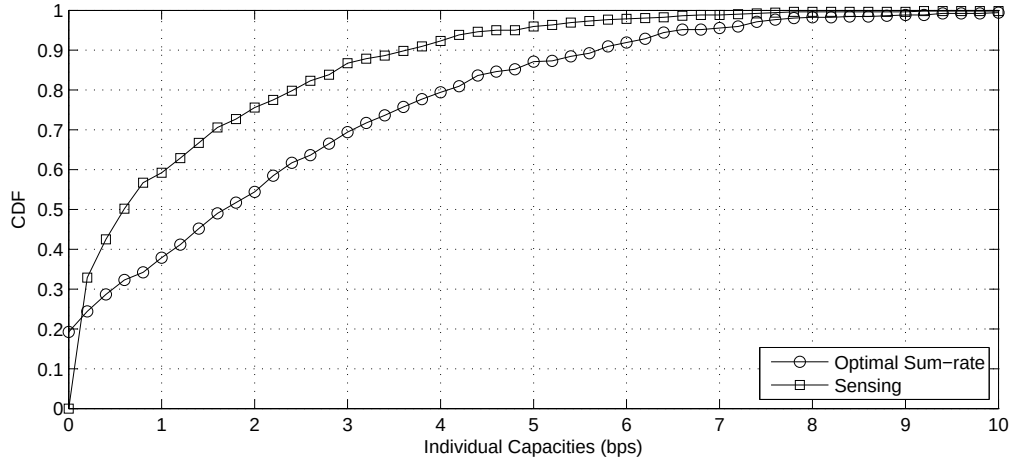

Figure 4.28: CDF of Individual Capacities of the Middle Two Users ($MUI = 1$, $SNR_{min} = 3dB$, $N = 6$, and $M = 2$)

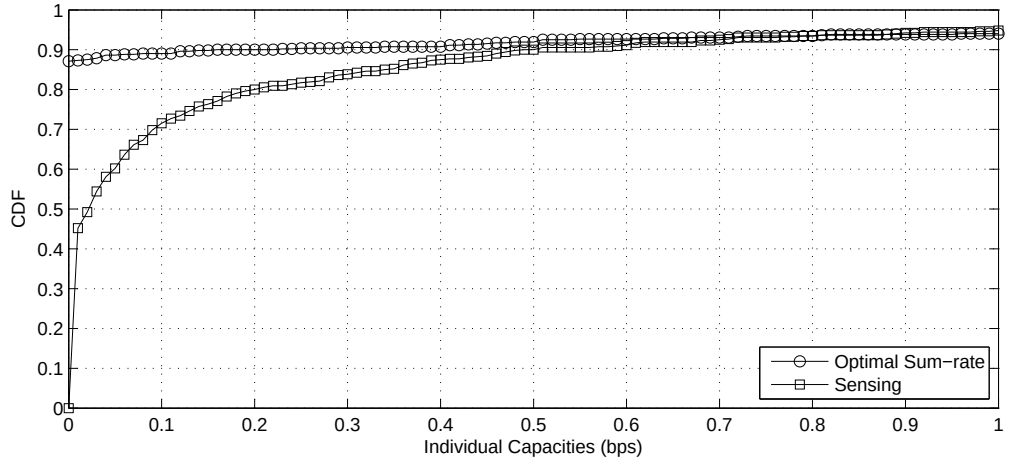

Figure 4.29: CDF of Individual Capacities of the Worst Two Users ($MUI = 1$, $SNR_{min} = 3dB$, $N = 6$, and $M = 2$)

performance while optimizing sum-rate with a Jain fairness goal of 0.5 lies in the middle between these two and optimizing sum-rate. While having a larger MUI causes the benefit of the collaboration to increase, it has little impact on the performance comparison of the techniques.

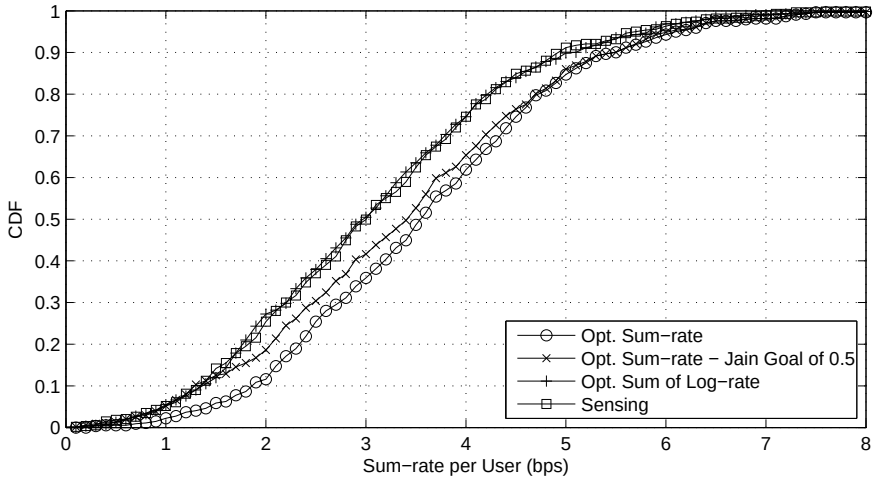

Figure 4.30: Sum-rate per User ($MUI = 1$, $SNR_{min} = 3dB$, $N = 6$, and $M = 2$)

In Figures 4.32 and 4.33, we plot the CDF of the Jain fairness assuming the same parameters ($MUI = 1$ for the first figure and $MUI = 5$ for the second figure). We note that because the Jain fairness is constrained to be higher than 0.5 when we optimize sum-rate with the Jain fairness goal, there is a jump in the CDF at a Jain fairness of 0.5. We note from the CDF that the Jain fairness optimization constraint fails only about 2% of the time. When Jain fairness is higher than 0.5, the optimizing sum-rate and optimizing sum-rate with the Jain fairness goal CDF curves are almost identical. We further note that sensing tends to have the worst Jain fairness while optimizing sum of log-rate tends to have the best. We note that, while the average Jain fairness for optimizing the sum-rate with a Jain fairness goal might be equal to or higher than the average Jain fairness for optimizing the log of sum-rate, it is hard to compare the impact of each technique without studying the outage probability plot. Therefore, we consider the outage probability plot below. We note that the fairness of optimizing sum-rate is slightly better than the fairness of sensing, and this matches a result shown in [75]. In this work, the cooperative approach yields both a slight capacity benefit and a slight fairness benefit over the selfish approach.

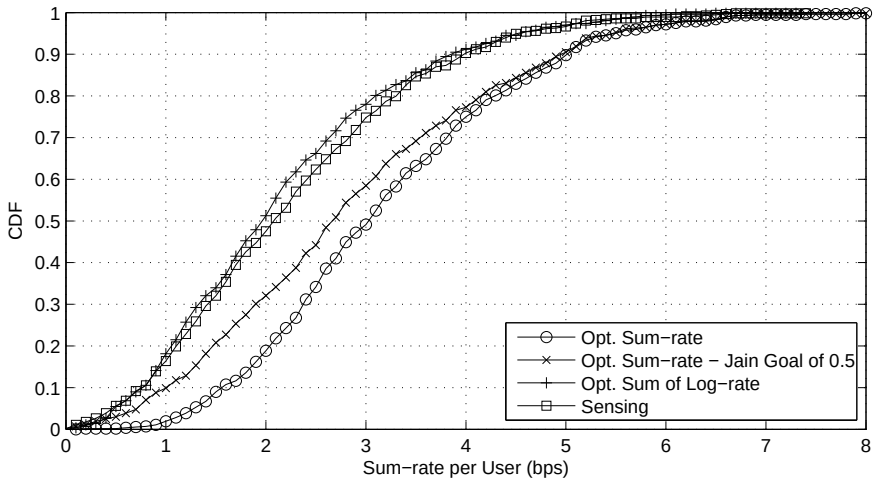

Figure 4.31: Sum-rate per User ($MUI = 5$, $SNR_{min} = 3dB$, $N = 6$, and $M = 2$)

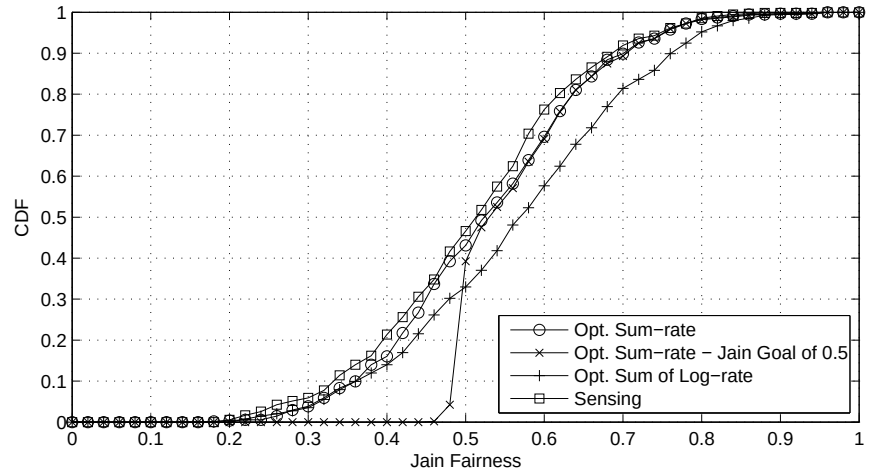

Figure 4.32: CDF of Jain Fairness ($MUI = 1$, $SNR_{min} = 3dB$, $N = 6$, and $M = 2$)

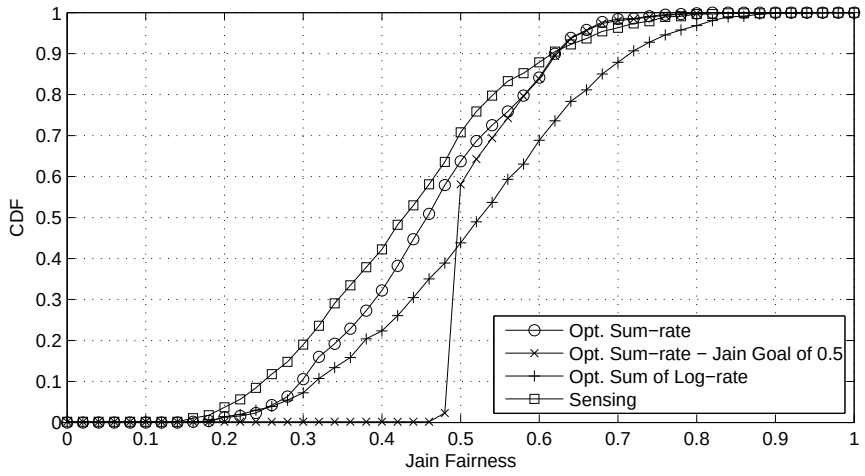

Figure 4.33: CDF of Jain Fairness ($MUI = 5$, $SNR_{min} = 3dB$, $N = 6$, and $M = 2$)

For the same input parameters, we plot the outage probability plots (CDF of the individual capacities) for $MUI = 1$ and $MUI = 5$ in Figures 4.34 and 4.35. In both cases, we note that optimizing sum-rate has the best capacities for the higher capacity users; however, many of the users receive no capacity (about 40% for $MUI = 5$). Optimizing sum-rate with the Jain fairness goal of 0.5 causes the highest capacity users to have slightly worse capacities, but only about 30% of users receive no capacity for $MUI = 5$. For both optimizing sum of log-rate and sensing, no users should receive no capacity (for sum of log-rate this would make the objective go to negative infinity and for sensing all users transmit full power so at least some minimal capacity is achieved). We can see that optimizing sum of log-rate tends to be slightly worse for the best capacity users than sensing, but it yields better outage capacities (i.e., for $MUI = 5$ the 30% outage capacity increases from about 0.15 for sensing to about 0.4 for optimizing sum of log-rate). So selection of which technique yields the best balance between fairness and capacity depends upon the goals of the system. If the system requires that all users have at least some minimum capacity, then optimizing sum of log-rate yields the best performance. If the system requires fairness but does not require that all users have some capacity, using a Jain fairness optimization goal while optimizing sum-rate tends to balance sum-rate and fairness. If no fairness constraints are considered, optimizing sum-rate tends to yield the best performance.

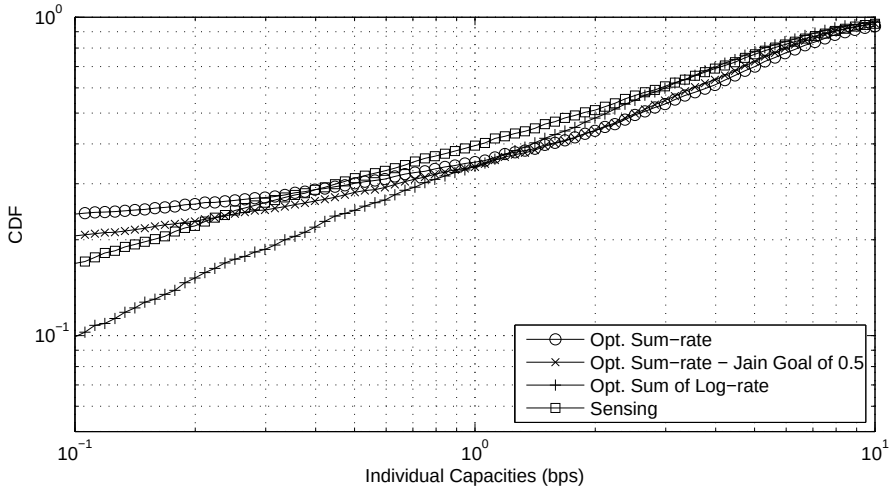

Figure 4.34: CDF of Individual Capacities ($MUI = 1$, $SNR_{min} = 3dB$, $N = 6$, and $M = 2$)

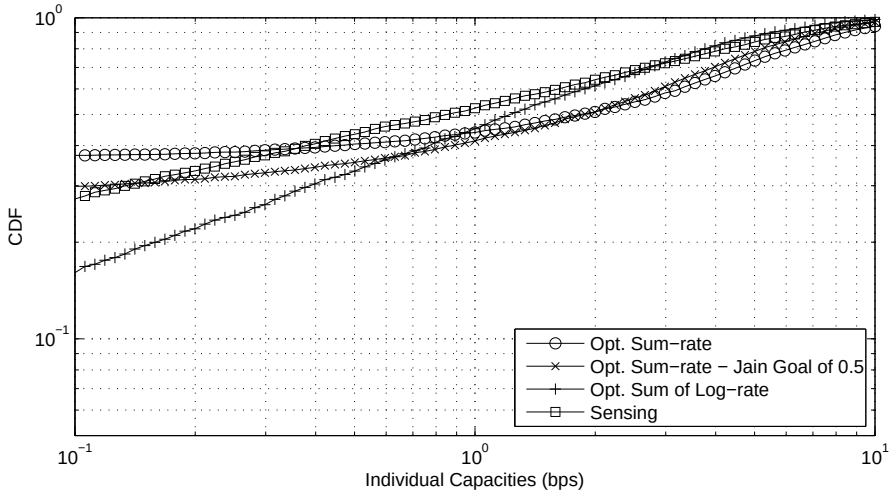

Figure 4.35: CDF of Individual Capacities ($MUI = 5$, $SNR_{min} = 3dB$, $N = 6$, and $M = 2$)

To illustrate the comparison of optimizing sum of log-rate with optimizing sum-rate with a Jain fairness goal, we compare the capacities of the best and worst users. In Figures 4.36, 4.37, and 4.38 we can see the CDF of the individual capacities of the best two users, middle two users, and worst two users, respectively. We can see that the addition of the Jain goal to the sum-rate optimization causes the performance of the best two users to be less, but it allows the middle and worst users to achieve slightly higher capacities. Optimizing the log of sum-rate gives the best performance to the worst users by far, but it yields the worst performance to the best users.

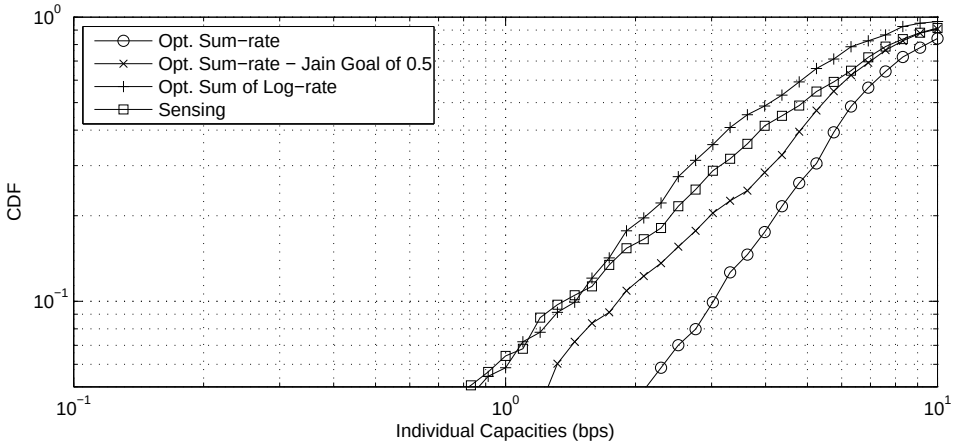

Figure 4.36: CDF of Individual Capacities of the Best 2 Users ($MUI = 5$, $SNR_{min} = 3dB$, $N = 6$, and $M = 2$)

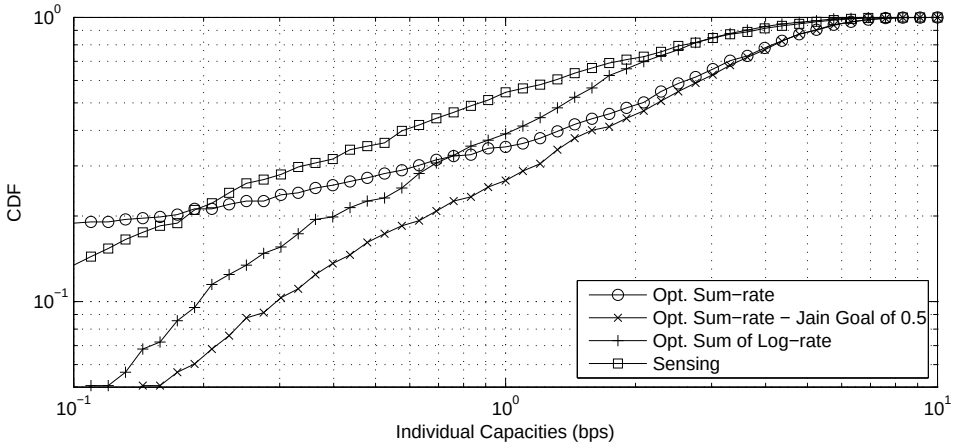

Figure 4.37: CDF of the Individual Capacities of the Middle 2 Users ($MUI = 5$, $SNR_{min} = 3dB$, $N = 6$, and $M = 2$)

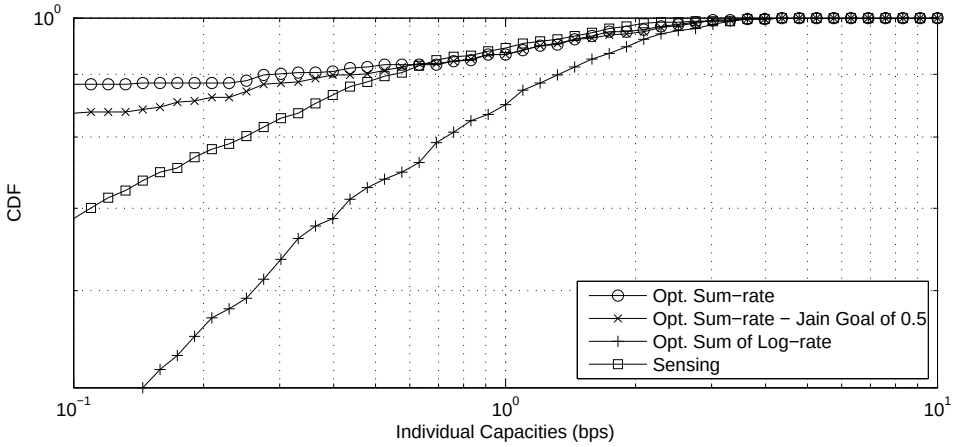

Figure 4.38: CDF of the Individual Capacities of the Worst 2 Users ($MUI = 5$, $SNR_{min} = 3dB$, $N = 6$, and $M = 2$)

To consider the impact of using different Jain fairness goals, we plot the CDF of sum-rate per user with no goal and with goals of 0.5 and 0.7 in Figure 4.39. We notice from the plot that as we increase the fairness of the system the sum-rate per user decreases. It is a common result in literature that maximizing capacity and fairness are conflicting goals (i.e., [30]). We can see the CDF of the individual user capacities in Figure 4.40. We can see that as we increase the Jain fairness goal fewer users tend to receive no capacity. However, increasing the Jain fairness causes the best users to receive worse performance.

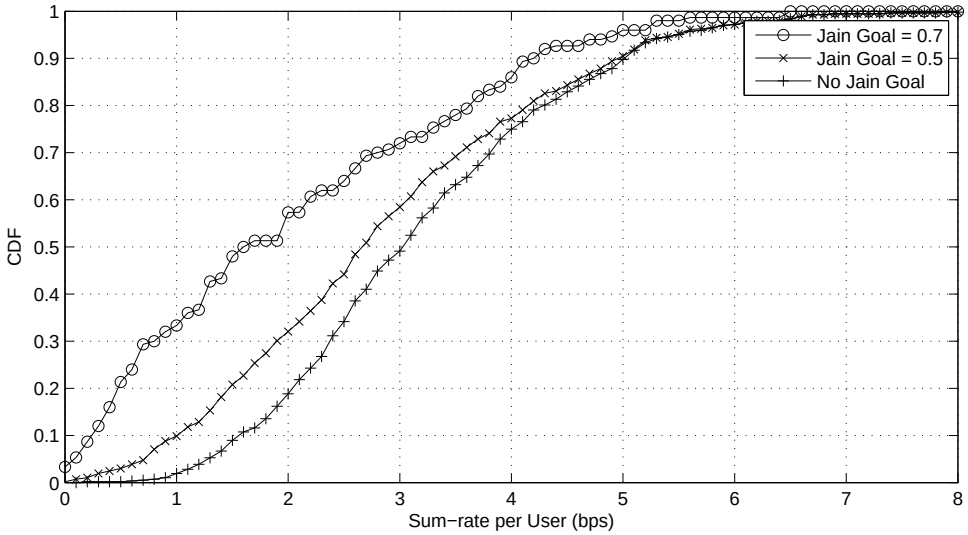

Figure 4.39: CDF of the Sum-rate per User ($MUI = 5$, $SNR_{min} = 3dB$, $N = 6$, and $M = 2$)

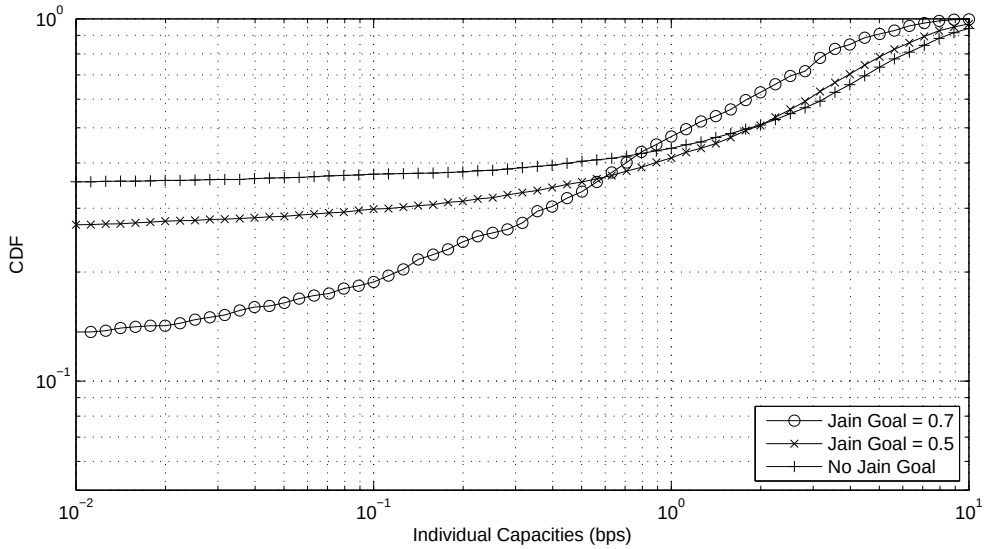

Figure 4.40: CDF of the Individual Capacities ($MUI = 5$, $SNR_{min} = 3dB$, $N = 6$, and $M = 2$)

4.5 Target-based Optimization

We now consider how optimization can improve the performance of the target-based sensing considered in Section 3.6. Remember that in that section we assumed that users had a selection between not transmitting or transmitting P_{max}. Because of this, we will specifically optimize the sum-rate per user (or equivalently the number of admitted users) assuming that all users can either transmit P_{max} or not transmit. We call this ability to either transmit maximum power or no power binary power control (BPC). In the next chapter, we will consider further the benefit of optimization which is allowed to use transmit powers P such that $0 \leq P \leq P_{max}$ (which we will call full power control (FPC)).

All simulations in this section assume $N = 6$, $M = 2$, $SNR_{min} = 3dB$ and $MUI = 1$. Recall from Section 3.6 that the information included in the sum-rate per user CDF and the Jain fairness CDF can be shown in the histogram of the number of active users. The histogram of the number of active users for both optimal (BPC) and target based sensing are plotted for a target rate of 1 in Figure 4.41 and a target rate of 3 in Figure 4.42. We immediately notice that for target rates of both 1 and 3 bps the performance of target based sensing is sub-optimal meaning there is a benefit to collaboration for target-based techniques. Similar to our result in Section 3.6 that for target based sensing the expected number of active users tends to decrease as the target rate increases, we notice that the expected number of active users for optimal BPC also decreases as the target rate increases. This is again due to fewer users being able to achieve the target rate. From the fraction of active users plot in Figure 4.43, we can see that increasing the target rate from 1 bps to 3 bps causes the performance of sensing techniques to be slightly closer to the optimal performance. As the target rate continues to increase, sensing should be able to perform closer to the optimal due to decreased interference caused by fewer users being active. In Figure 4.44, we plot the impact of changing the SNR_{min} to 9dB. We can see that while the fraction of active users increases, the benefit of collaborative approaches remains about the same. In Figure 4.45, we plot the fraction of active users if we increase the MUI to 5. As we increase the MUI, the fraction of active users decreases due to the increasing interference; however, the additional fraction of active users that we obtain by using optimization instead of sensing approaches remains nearly constant at slightly less than 10%.

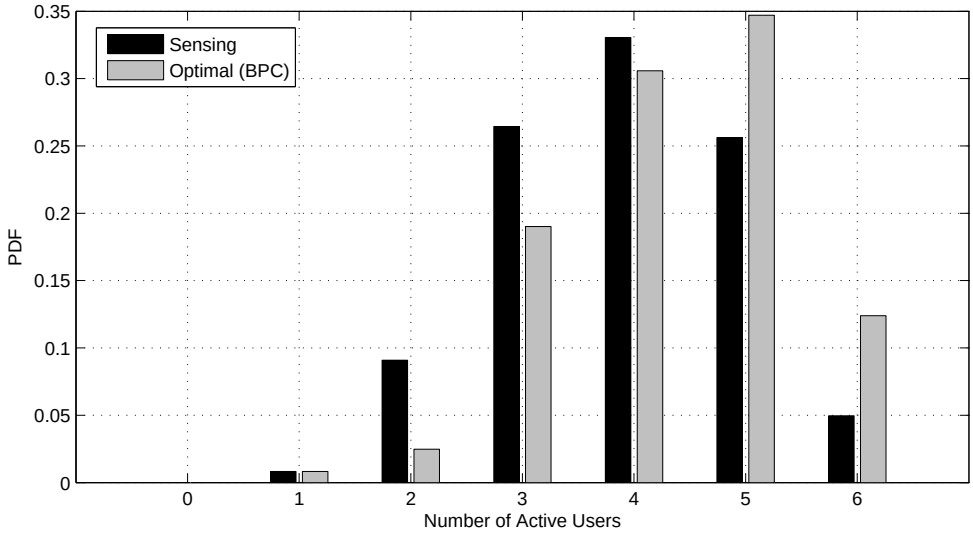

Figure 4.41: Histogram of Number of Active Users for a Target Rate of 1 bps ($MUI = 1$, $SNR_{min} = 3dB$, $N = 6$, and $M = 2$)

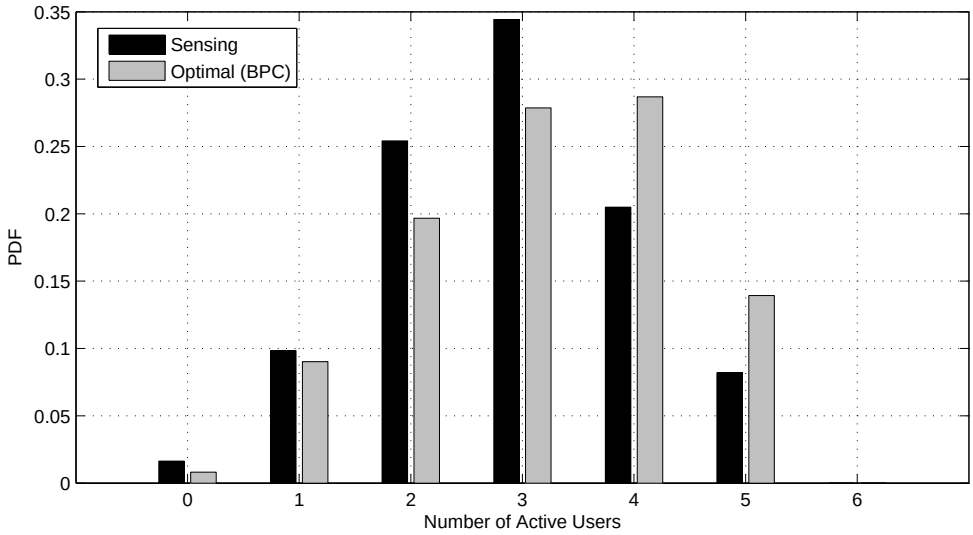

Figure 4.42: Histogram of Number of Active Users for a Target Rate of 3 bps ($MUI = 1$, $SNR_{min} = 3dB$, $N = 6$, and $M = 2$)

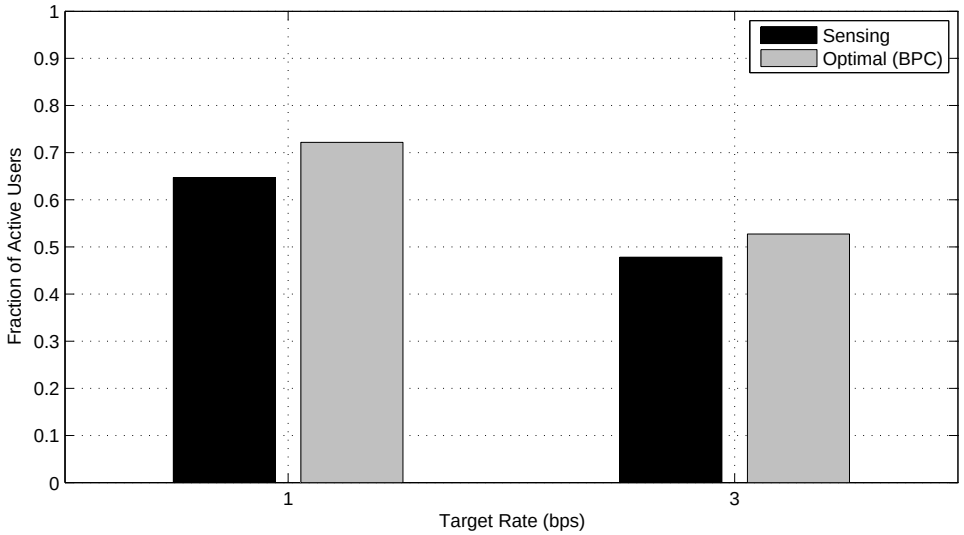

Figure 4.43: Fraction of Active Users for Target Rates of 1 bps and 3 bps ($MUI = 1$, $SNR_{min} = 3dB$, $N = 6$, and $M = 2$)

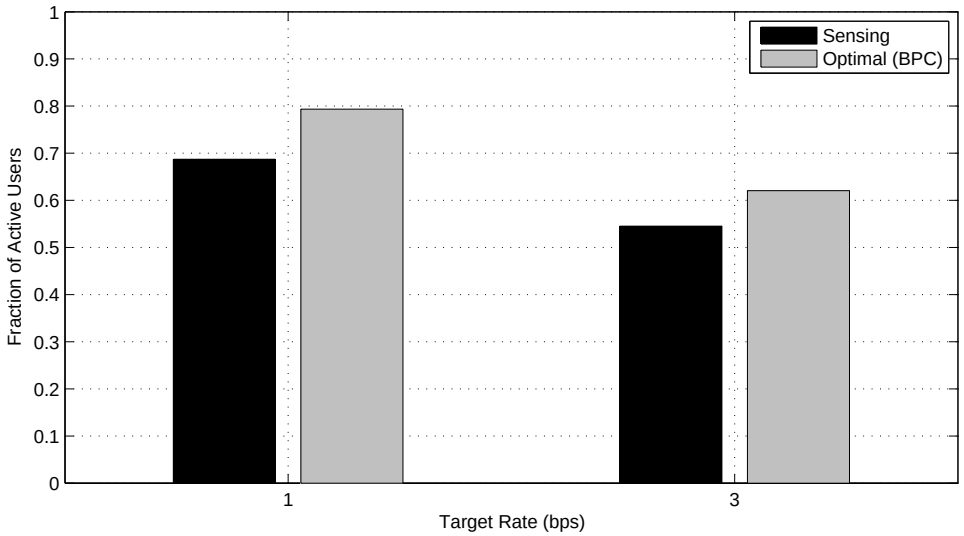

Figure 4.44: Fraction of Active Users for Target Rates of 1 bps and 3 bps ($MUI = 1$, $SNR_{min} = 9dB$, $N = 6$, and $M = 2$)

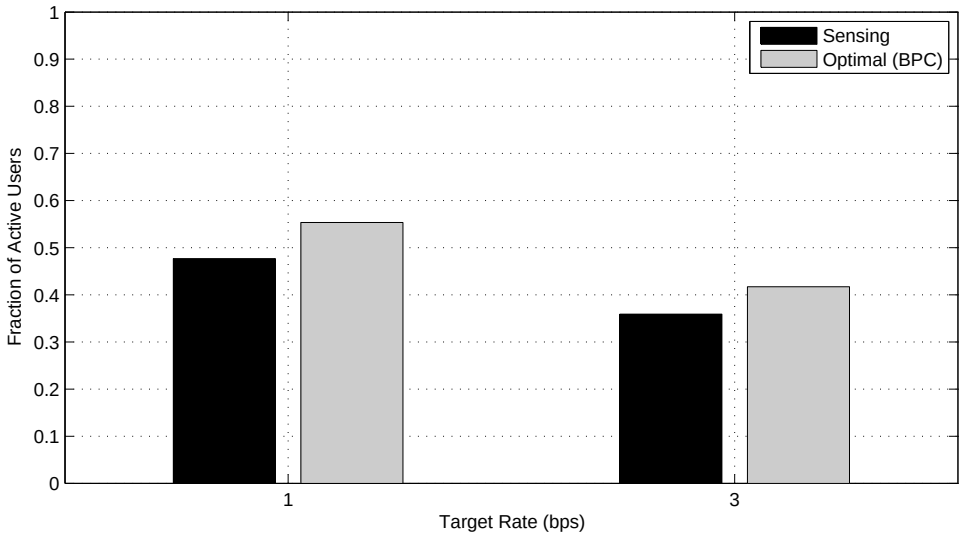

Figure 4.45: Fraction of Active Users for Target Rates of 1 bps and 3 bps ($MUI = 5$, $SNR_{min} = 3dB$, $N = 6$, and $M = 2$)

4.6 Conclusions

We began the chapter by introducing the optimization problem that is used to model full collaboration and verified that, by using three initial solutions with the Branch and Bound optimization technique, we obtain a result that is on average within 1% of the global optimum. We concluded the section by presenting simulation results and making conclusions about the impact of several parameters on the benefit of collaboration. These conclusions include:

- The capacity of our optimal sum-rate approach closely follows the trend introduced by Gupta and Kumar in [1].
- Collaboration in severe interference scenarios can result in gains as large as 100%; however, gains decrease as the number of channels increases, as the density of users decreases, and less than full collaboration is achieved. More realistic gains are in the range of 20-30%.
- Increased shadowing causes the collaboration gain to increase.
- Slow (predictable) fading has a large capacity benefit over fast (non-predictable) fading. However, since both impact sensing and optimization evenly, there is relatively little impact on the collaboration gain whether slow of fast fading is assumed.
- Increasing the number of channels increases the multi-channel diversity, but the capacity benefit follows the law of diminishing returns. As the number of channels increases and users are spread amongst more channels, the expected interference power decreases causing the collaboration gain to decrease.
- Increasing the number of users decreases the average sum-rate per user. It simultaneously increases sum-rate by increasing multi-user diversity and decreases sum-rate per user since it adds an additional user. As the number of users increases, the expected interference power increases causing the collaboration gain to increase.
- When the noise is dominant over interference, there is very little benefit to collaboration; however, when the interference is dominant over the noise, there is a relatively high benefit to collaboration. This results because collaboration can reduce cross-user interference.
- Adding outside interference quickly degrades the benefit of collaboration (even at low loading values). Adding outside interference is nearly equivalent to raising the noise floor.
- Using bi-directional links without reoptimizing each time a user switches transmission direction has little impact on sum-rate or the benefit of collaboration (benefit of collaboration decreased by a maximum of around 5%).

- Partial collaboration decreases the collaboration gain at a rate that is nearly linear with the number of groups of collaborating users.
- By optimizing sum-rate, we tend to increase fairness and capacity over using sensing simultaneously. This matches the results found in [75] which show the collaborative approach to achieve better capacity and fairness than the non-collaborative approach.
- Optimizing capacity tends to be obtained at the expense of fairness (this is also noted in [30]); however, collaboration can be used to increase fairness by using the proper objective function.
- While fairness is defined as evenly distributing capacity amongst all users, we showed two different optimization techniques that yielded results that were 'fair' in their own means. Optimizing sum-rate with a Jain fairness goal increased the number of included users but maintained a higher capacity than optimizing sum of log-rate for the same average Jain fairness. Optimizing sum of log-rate required all users receive some capacity at the cost of degrading performance to the best users.
- Collaboration in a target-based system tends to give benefit over target-based sensing; however, these gains are relatively small ($< 10\%$ increase in the number of users achieving the target rate).

Chapter 5

Value of Power Control / Multiband

In the previous chapter, we introduced the optimization problem used to model full collaboration and determined the impact of system parameters on the benefit of collaboration. In this chapter, we will study the impact of allowing users to use differing degrees of freedom, namely, different levels of power control and multiband approaches. We begin the chapter by introducing the different levels of power control considered in this work. Through simulation results, we determine the impact of different levels of power control on the benefit of collaboration. For the specific case of two users and one channel, we prove that using binary power control and using full power control achieve equivalent results. We then consider whether using multiband approaches (such as multi-channel (MC) and spread spectrum (SS)) has any impact on the benefit of collaboration through simulation. We consider some specific scenarios for the case with two users and two channels and prove that for these scenarios spread spectrum approaches do not yield any benefit over single-channel approaches. We consider some example multi-channel optimization results to explain the benefit of using multi-channel techniques. We conclude by reviewing the impact of multiband approaches on the benefit of collaboration.

5.1 Levels of Power Control

We begin by defining the different types of power control that can be implemented within the optimization routine. These include full power control, binary power control, and no power control. With full power control, we assume that each Tx node can transmit any power level P in the range $0 \leq P \leq P_{max}$. (Note that all previous simulations assume full power control.) Binary power control (also termed on/off power control) is defined as the power control where each Tx node must either transmit 0 or P_{max} and thus is restrained to a binary selection of whether to transmit full power or no power. When no power control is used, the Tx node transmits P_{max} and has no control over the power value.

Since binary power control is the same as no power control but has additional freedom, the performance of binary power control should be equal to or better than that of no power control. Similarly, the performance of full power control should be equal to or better than that of both binary power control and no power control. By comparing the results of using different types of power control and sensing, we can determine what portion of the collaboration gains can be attributed to channel selection, to being able to turn users off, and to continuous power control. For example, by comparing the optimal no power control case to sensing, we have completely removed power control and are only considering the benefit of optimal channel selection. We will study these relationships in more detail in Chapter 6.

5.2 Simulation Results for Power Control

In this section, we simulate full collaboration utilizing different types of power control for $MUI = 5$ and $SNR_{min} = 3dB$. In the figures, we denote full power control by the acronym FPC, binary power control by the acronym BPC, and no power control by the acronym NPC. The CDF of the sum-rate per user for $N = 6$ and $M = 1$ can be seen in Figure 5.1. We can see that the performance of binary power control and full power control are nearly identical. Furthermore, the performance of no power control optimization and sensing are identical. Because the difference between NPC and sensing is channel selection and we only have a single channel, no benefit exists for NPC over sensing. For this scenario, we can see that the largest benefit of collaboration is binary power control or being able to turn users off. The CDF of the sum-rate per user for $N = 6$ and $M = 3$ can be seen in Figure 5.2. Again the performance of BPC and FPC are nearly identical; however, NPC is now also nearly identical with full power control. By increasing the number of channels to 3, the benefit of channel selection increases while the benefit of turning users off decreases.

We expect the benefit of channel selection to exist only when the number of channels is larger than 1 and to decrease as the number of channels increases past 2. As explained previously, no benefit to channel selection can exist if only a single channel is present. As the number of channels increases, the benefit of channel selection decreases because the interference per channel decreases as users spread amongst more channels causing the interference avoidance (the main benefit of collaboration) to become less significant. When we have very few channels and many users, a benefit of turning users off exists, but this benefit quickly disappears when all users can be accommodated without causing harmful interference to each other (i.e., when the number of users decreases or the number of channels increases). Full power control tends to have very little benefit over binary power control when optimizing sum-rate. In Section 5.3, we will prove that, when 2 users share a single channel, optimizing sum-rate with full power control and optimizing sum-rate with binary power control yield equivalent results.

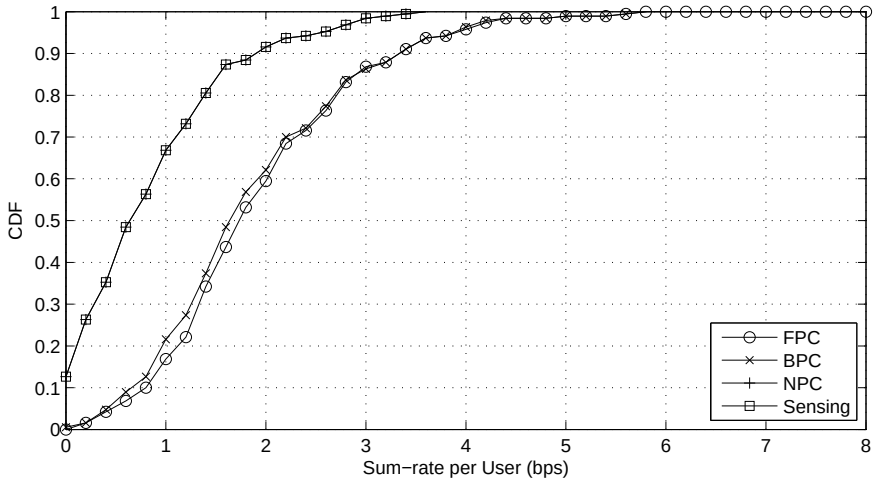

Figure 5.1: CDF of the Sum-rate per User with 6 Users, 1 Channel

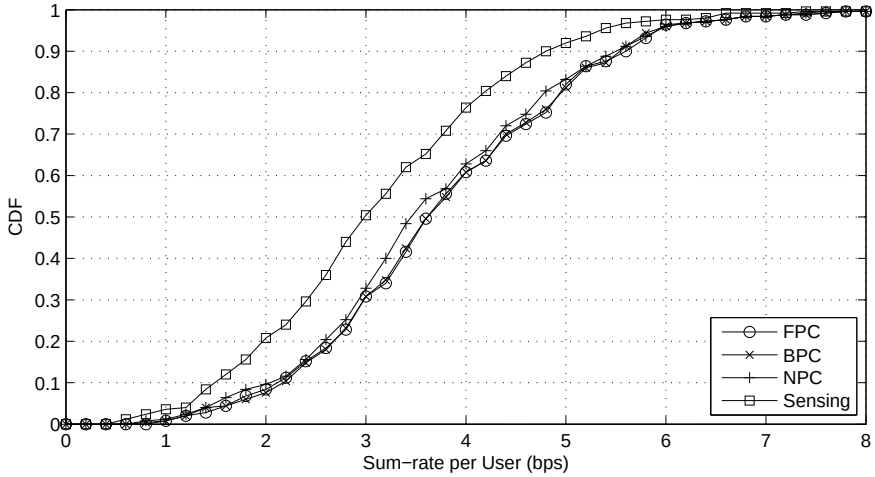

Figure 5.2: CDF of the Sum-rate per User with 6 Users, 3 Channels

We now want to consider the impact of different types of power control on fairness. To do this, we plot the CDF of the Jain fairness for the $N = 6$, $M = 1$ case in Figure 5.3. Like in the sum-rate per user CDF, the binary power control and full power control curves are nearly identical, and the no power control and sensing curves are identical. We can see that, while the average Jain fairness for all methods are nearly the same, there are times when the Jain fairness of binary power control and full power control is better and times when the Jain fairness of no power control and sensing is better. The CDF of the Jain fairness for $N = 6$, $M = 3$ can be seen in Figure 5.4. All three optimization approaches are nearly identical and tend to yield better fairness than the sensing approach. The Jain fairness of the optimization results are all similar because we expect the optimization results using all three levels of power control to be similar (since the sum-rate per user plots are similar, we assume that using NPC can often achieve the same optimal result that FPC achieves).

Two different effects impact the fairness of the optimization results when full power control is used. The optimization affects the fairness negatively by prioritizing the users with the best channels (neglecting the users with the worst channels) but it also affects fairness positively by causing the best users to all have similar performance. The prioritizing of the users with the best channels can be seen in the individual capacity CDF for $N = 6$, $M = 1$ in Figure 5.5. In this figure, the FPC and BPC curves are nearly flat at low capacities at around the 60% line signifying that 60% of the users receive no capacity. If we examine the histogram of the individual capacities in Figure 5.6, we can see that by sacrificing the lower capacity users FPC and BPC optimizations give the best users similar performance. We can see from the CDF of $N = 6$ and $M = 3$ in Figure 5.7 that about 20% of users receive no capacity. In Figure 5.8, we plot the histogram of individual user capacities for $N = 6$ and $M = 3$. Unlike when we had a single channel, the peak in individual capacities between 1 bps and 10 bps is only slightly higher for the optimization techniques than it is for the sensing techniques. We note that for both Figures 5.6 and 5.8 there is an impulse which cannot be seen at zero for BPC and FPC because these two methods allow users to be turned off.

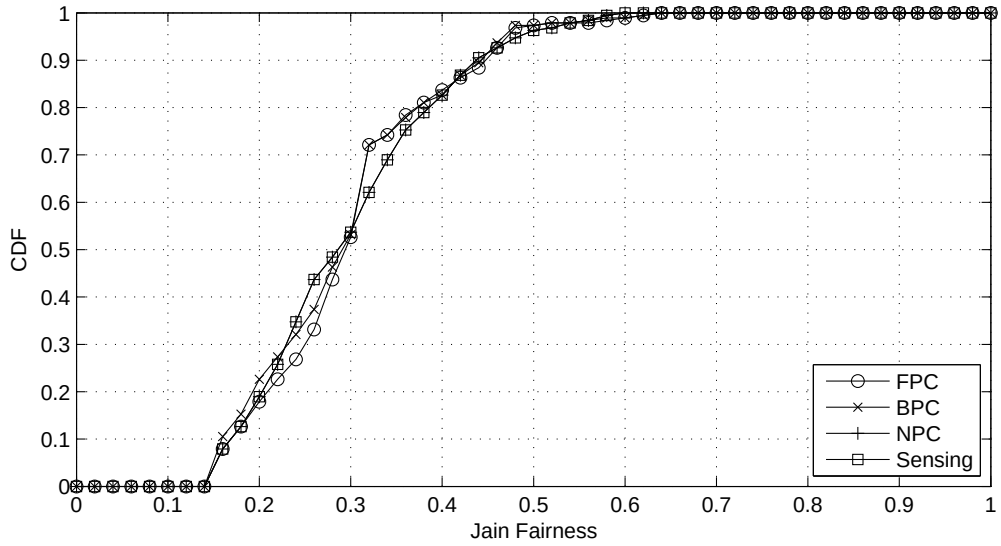

Figure 5.3: CDF of the Jain Fairness with 6 Users, 1 Channel

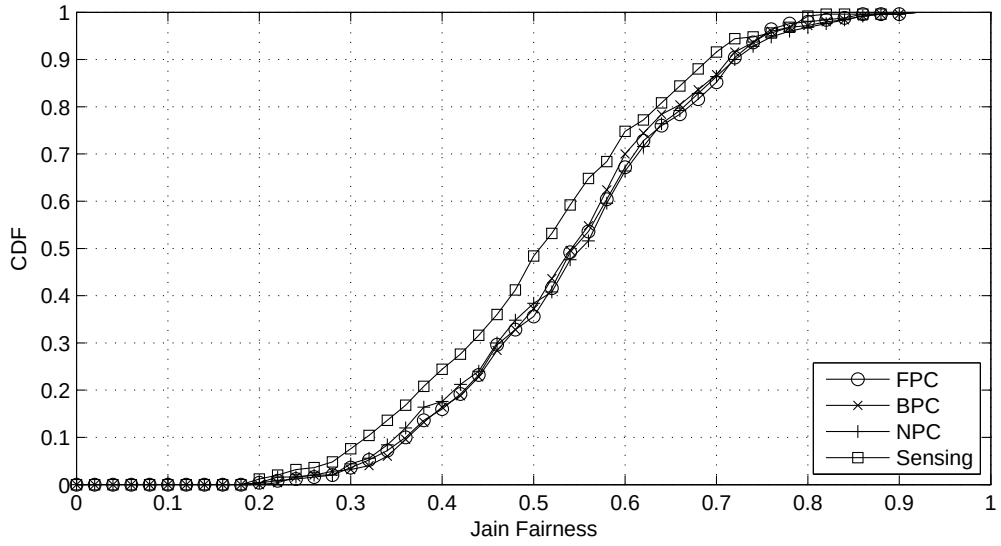

Figure 5.4: CDF of the Jain Fairness with 6 Users, 3 Channels

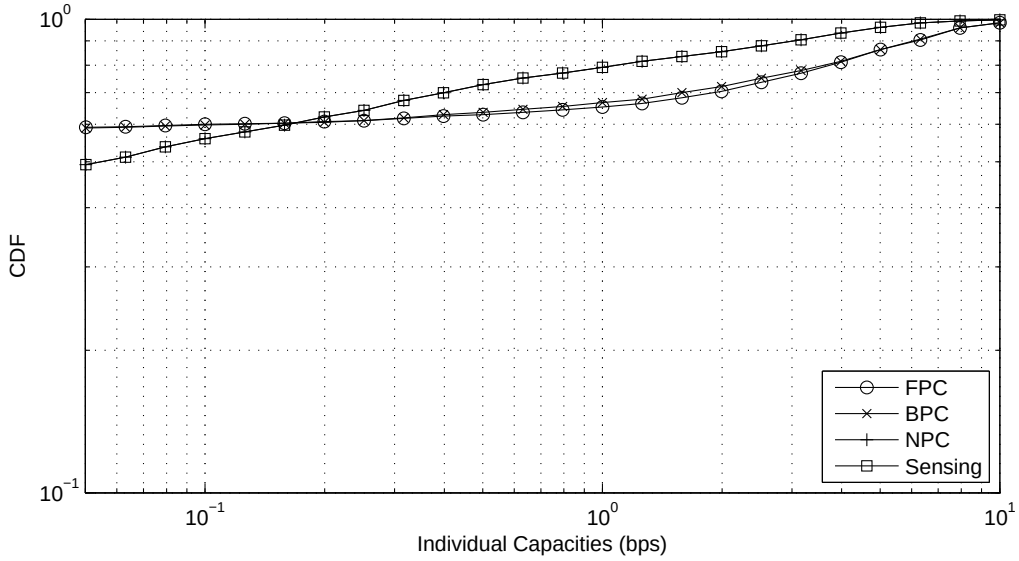

Figure 5.5: CDF of the Individual User Capacities with 6 Users, 1 Channel

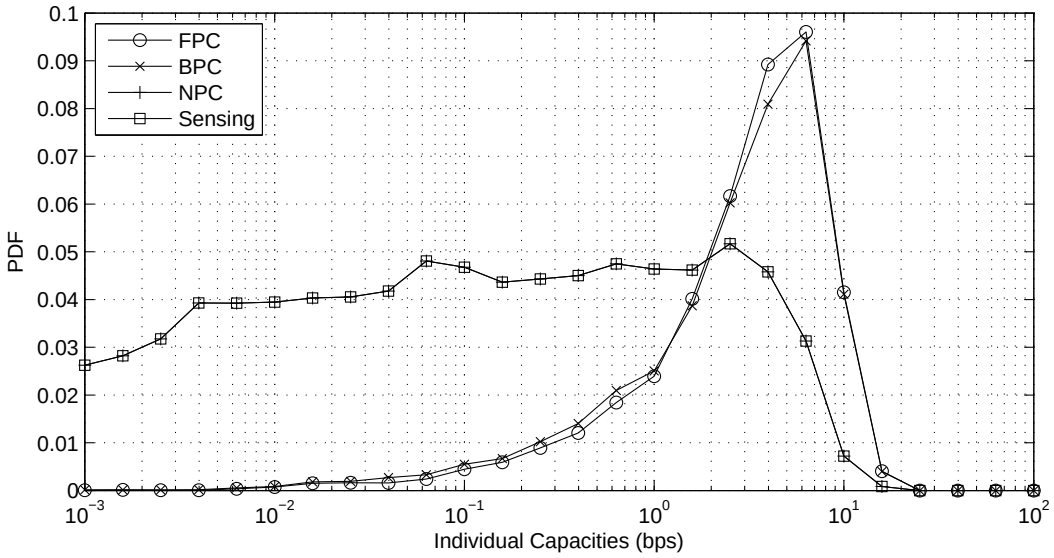

Figure 5.6: Histogram of the Individual User Capacities with 6 Users, 1 Channel

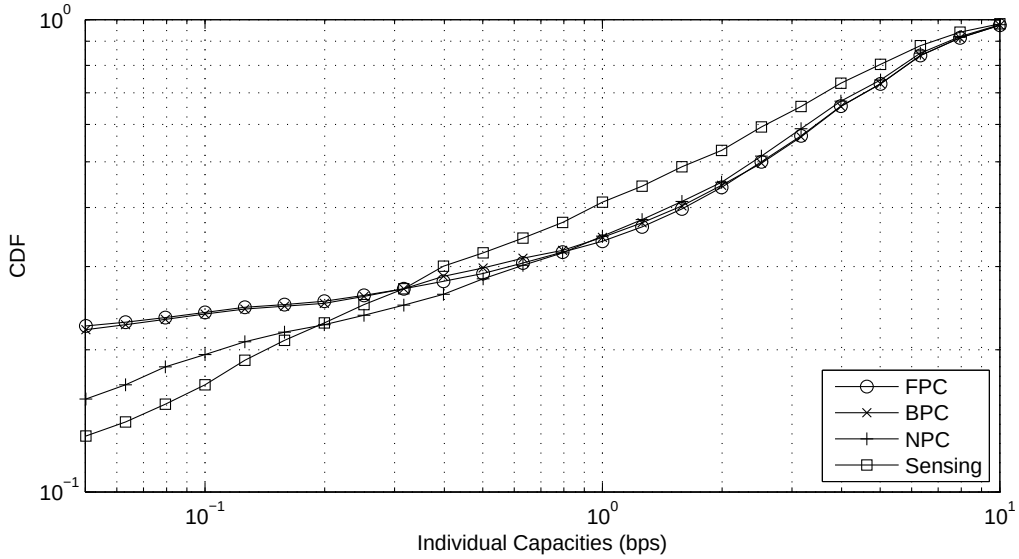

Figure 5.7: CDF of the Individual User Capacities with 6 Users, 3 Channels

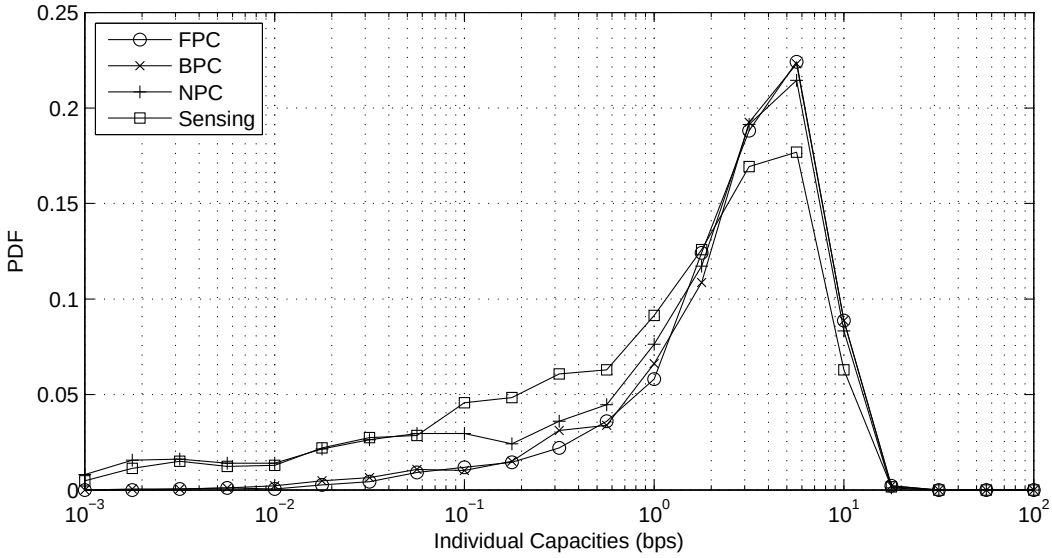

Figure 5.8: Histogram of the Individual User Capacities with 6 Users, 3 Channels

For the sum-rate optimization problem, we discovered that little benefit exists for full power control over binary power control. Now we consider whether full power control has any benefit over binary power control for the sum of log-rate optimization problem. We note that because the sum of log-rate optimization problem requires that all users have non-zero capacities, no power control and binary power control are equivalent (if all users are required to have a non-zero capacity, no user can transmit zero power). For the case of $N = 6$ and $M = 2$, we plot the CDF of the sum-rate per user in Figure 5.9 and the CDF of the Jain fairness in Figure 5.10. We note that the sum-rate per user of binary or no power control is better than than that of full power control, but the fairness of full power control is better than that of binary or no power control. This is true because, without power control, the users with the best channels cannot decrease their transmit powers to allow the users with the worst channels to achieve better performance. While no power control causes the users with the best channels to have better performance (which has a positive impact on sum-rate) it causes the users with the worst channels to have worse performance (which has a negative impact on fairness). Thus full power control does provide some benefit when fairness is a major concern.

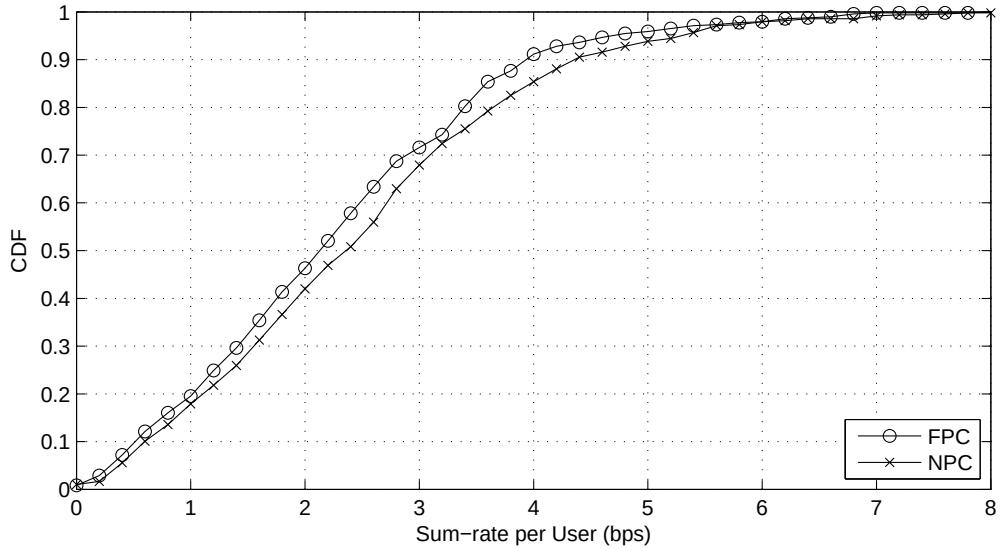

Figure 5.9: CDF of the Sum-rate per User Optimizing Proportional Fairness (6 Users, 2 Channels)

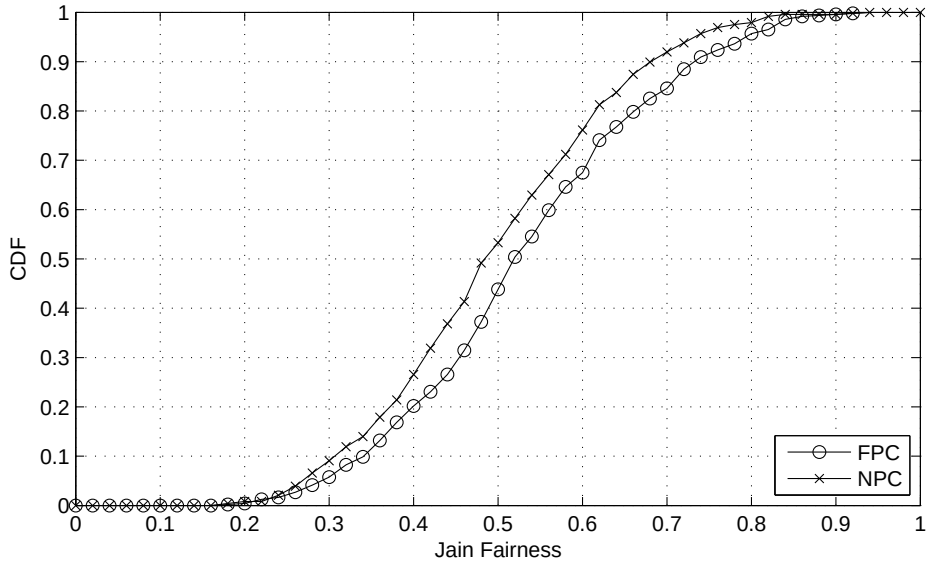

Figure 5.10: CDF of the Jain Fairness Optimizing Proportional Fairness (6 Users, 2 Channels)

The optimization problem that many cellular systems use is maximizing the number of users allowed to transmit given that all transmitting users must achieve a fixed capacity value (which we call the target rate). In many works, full power control is emphasized as being an important part of target based cellular systems with cochannel interference (i.e., [94]). In [95], Zander explains the benefit of full power control by saying that the power level must be high enough for the users to achieve the target rate but it should simultaneously minimize interference to other users. This explanation is the same as the explanation for the sum of log-rate optimization where we said that the best user's power levels must be high enough for these users to achieve good capacities, but it should not be so high that it causes excessive harmful interference to the worst users.

We now consider the optimization problem where we maximize the number of users active on a set of channels given that all active users must achieve a target capacity. To enforce this, we implement a global minimum bps limit and global maximum bps limit which are equal to the target rate. This means that each user in the optimization either achieves the target capacity or has zero capacity.

Recall from Section 3.6 that instead of plotting the CDF of sum-rate per user and the CDF of Jain fairness, we plot the histogram of the number of active users. In Figures 5.11 and 5.12, we plot the histogram of the number of active users with $N = 6$ and $M = 2$ with target rates of 1 and 3. We notice immediately that full power control has a noticeable benefit over binary power control as expected. In comparing the fraction of active users of BPC and FPC for a MUI of 1 in 5.13 and a MUI of 5 in 5.14, we notice that the benefit of full power control over binary power control seems to be slightly higher when the target rate is lower. This is because increasing the target rate decreases the number of users able to contend for the channel. As contention decreases, the benefit of interference minimization decreases making the benefit of FPC decrease. We can also see that as the MUI increases the benefit of using FPC over BPC increases slightly due to the increase in expected interference.

We note that the optimization problem with FPC is a difficult one (Zander also arrived at this conclusion in [95]), so our FPC results are a lower bound on the optimal performance. For BPC, our results are guaranteed to be optimal because we used a brute force approach in which we tested every possible combination of users transmitting on channels. Combining these facts, the gains we show here for using FPC over BPC are a lower bound on the true gains achievable.

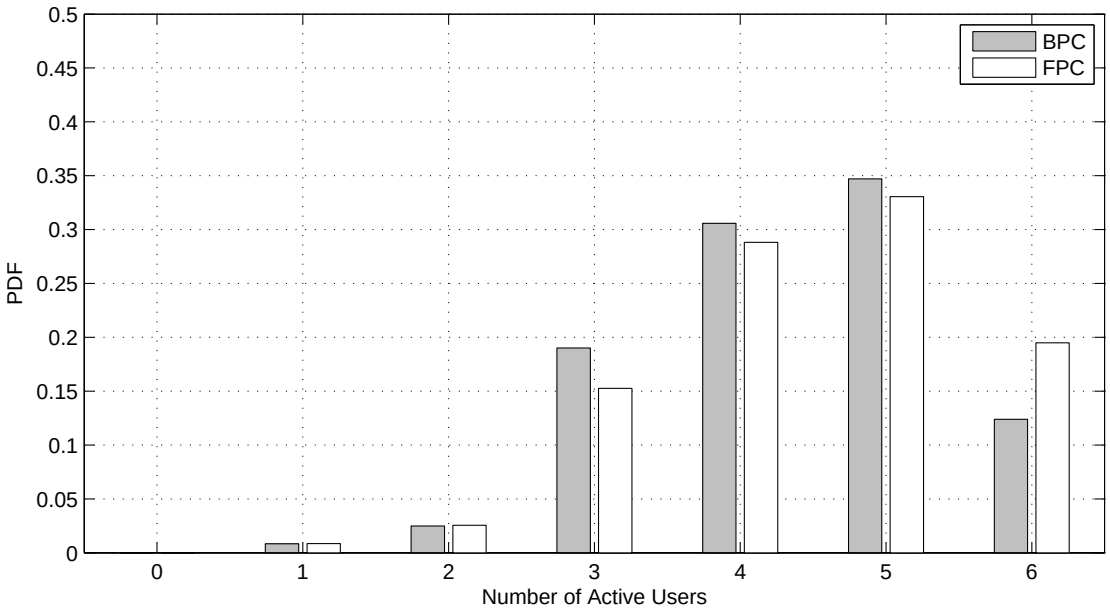

Figure 5.11: Histogram of Number of Active Users with a Target Rate of 1 bps (6 Users, 2 Channels, $SNR_{min} = 3dB$, $MUI = 1$)

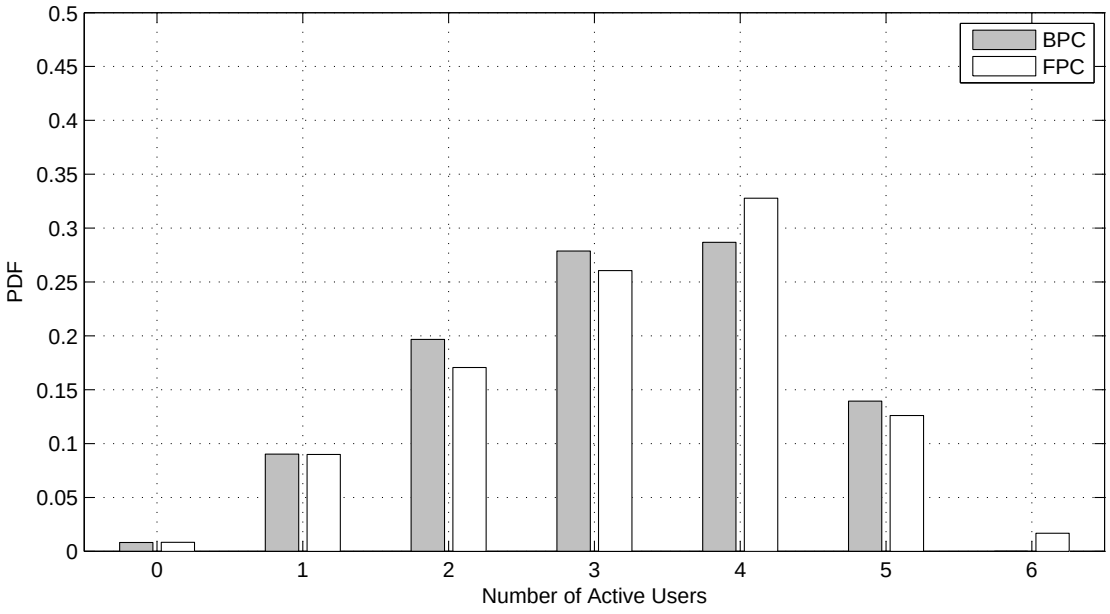

Figure 5.12: Histogram of Number of Active Users with a Target Rate of 3 bps (6 Users, 2 Channels, $SNR_{min} = 3dB$, $MUI = 1$)

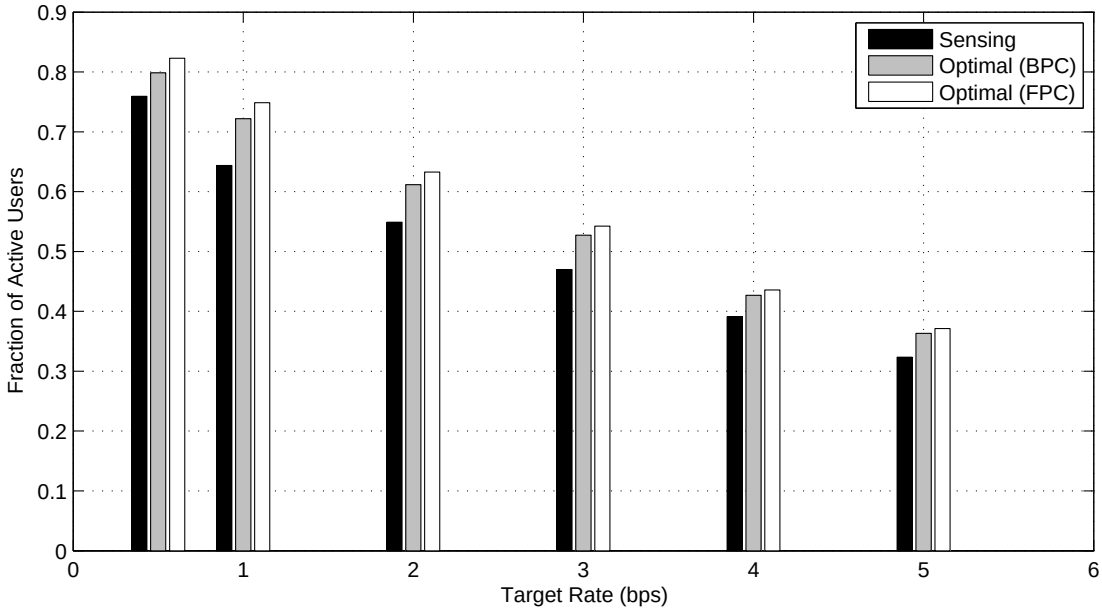

Figure 5.13: Fraction of Active Users for Varying Target Rates (6 Users, 2 Channels, $SNR_{min} = 3dB$, $MUI = 1$)

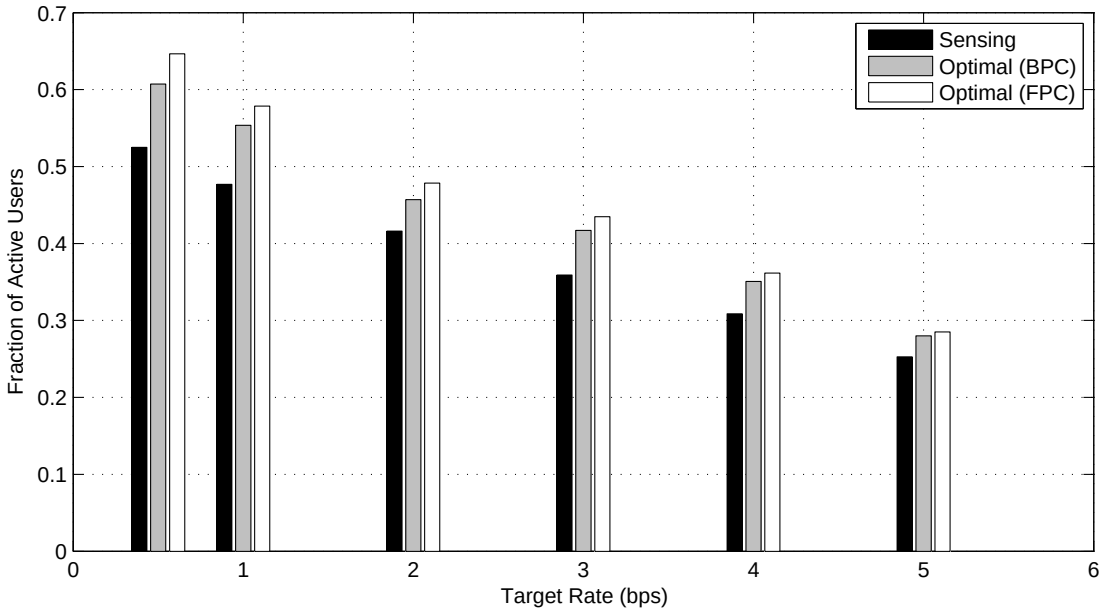

Figure 5.14: Fraction of Active Users for Varying Target Rates (6 Users, 2 Channels, $SNR_{min} = 3dB$, $MUI = 5$)

5.3 Equivalence of FPC and BPC: Proof for the Two-User Case

In the previous section, we showed through simulation that, when optimizing sum-rate, using binary power control and using full power control yield nearly equivalent results. These results mean that in the strategy that optimizes sum-rate using full power control, most Tx nodes either transmit 0 or P_{max}. In this section, we demonstrate that for the two-user, one channel case this is indeed the optimal result. In other words, the optimal solution is that both users transmit at either P_{min} or P_{max} and thus binary power control is identical to full power control.
For each of the two users, we can represent the $SINR$ values by

$$\begin{aligned} SINR_1 &= \frac{P_1 \cdot g_{1,1}}{P_2 \cdot g_{2,1} + \nu_1} \\ SINR_2 &= \frac{P_2 \cdot g_{2,2}}{P_1 \cdot g_{1,2} + \nu_2} \end{aligned} \tag{5.1}$$

The capacity of the individual users is then given in (5.2) and the sum-rate is given by (5.3).

$$\begin{aligned} C_1 &= log_2\left(1 + SINR_1\right) \\ C_2 &= log_2\left(1 + SINR_2\right) \end{aligned} \tag{5.2}$$

$$\begin{aligned} C_{Total} &= C_1 + C_2 \\ &= log_2\left(1 + SINR_1\right) + log_2\left(1 + SINR_2\right) \\ &= log_2\left(\left(1 + SINR_1\right) \cdot \left(1 + SINR_2\right)\right) \end{aligned} \tag{5.3}$$

Lemma 1 *For the case of a single frequency band, in order to optimize sum-rate for the two-user case, either P_1 or P_2 must equal P_{max}.*

Proof. Let us assume that this statement is not true and $P_1, P_2 < P_{max}$ while we are achieving optimal sum-rate. We know that if we can increase both $SINR_1$ and $SINR_2$ simultaneously, we will increase both user's individual capacities and we will increase the sum-rate. Therefore if we have achieved the optimal sum-rate, it must be impossible to increase both $SINR_1$ and $SINR_2$ simultaneously. Let us increase P_1 and P_2 by multiplying both by $1 + \Delta$ where $\Delta > 0$, $(1 + \Delta) \cdot P_1 < P_{max}$ and $(1 + \Delta) \cdot P_2 < P_{max}$. Note that in (5.1) this has an equivalent effect of multiplying the noise terms ν_i by $\frac{1}{1+\Delta}$. By effectively decreasing the noise power terms and thus decreasing the denominator, we increase both $SINR_1$ and $SINR_2$ simultaneously. This means that our original statement that we had achieved optimal sum-rate is false, so either $P_1 = P_{max}$ and/or $P_2 = P_{max}$ to reach optimal sum-rate. ■

Lemma 2 *Given the validity of Lemma 1, the optimal sum-rate is achieved by one of three points:* $\{P_{max}, P_{min}\}$, $\{P_{min}, P_{max}\}$, *or* $\{P_{max}, P_{max}\}$.

Proof. Without loss of generality, let us assume that $P_2 = P_{max}$ and $0 \leq P_1 \leq P_{max}$ ($P_{min} = 0$). We must then show that P_1 is either equal to 0 or P_{max} when optimized over sum-rate. If no maximum of sum-rate is reached while P_1 in varied from 0 to P_{max}, then we have proven that P_1 is equal to either 0 or P_{max}. This is because, if no maximum exists between the two points, then we can continually vary P_1 in the direction that increases sum-rate until we reach either 0 or P_{max}. Since a maximum would mean that the derivative of sum-rate with respect to P_1 is equal to zero, we take the derivative in (5.4).

$$\begin{aligned}\frac{\partial C_{Total}}{\partial P_1} &= ln\left(2\right) \cdot \frac{1}{\left(\left(1+SINR_1\right)\cdot\left(1+SINR_2\right)\right)} \cdot ... \\ &\left(\left(\left(1+SINR_1\right)\cdot\frac{\partial SINR_2}{\partial P_1}\right)+\left(\frac{\partial SINR_1}{\partial P_1}\cdot\left(1+SINR_2\right)\right)\right) \end{aligned} \tag{5.4}$$

Since we care about where this derivative is equal to zero and we know that $ln\left(2\right) \neq 0$ and $\frac{1}{\left(\left(1+SINR_1\right)\cdot\left(1+SINR_2\right)\right)} \neq 0$, the maxima/minima occur when

$$\left(\left(\left(1+SINR_1\right)\cdot\frac{\partial SINR_2}{\partial P_1}\right)+\left(\frac{\partial SINR_1}{\partial P_1}\cdot\left(1+SINR_2\right)\right)\right) = 0 \tag{5.5}$$

The partial derivatives of the $SINR$ values with respect to P_1 can be shown to be

$$\frac{\partial SINR_2}{\partial P_1} = \frac{-g_{1,2}\cdot P_2 \cdot g_{2,2}}{\left(P_1\cdot g_{1,2}+\nu_2\right)^2} \tag{5.6}$$

$$\frac{\partial SINR_1}{\partial P_1} = \frac{g_{1,1}}{P_2\cdot g_{2,1}+\nu_1} \tag{5.7}$$

By inserting these values into (5.5) and simplifying,

$$\left(\left(\left(1+SINR_1\right)\cdot\frac{-g_{1,2}\cdot P_2\cdot g_{2,2}}{\left(P_1\cdot g_{1,2}+\nu_2\right)^2}\right)+\left(\frac{g_{1,1}}{P_2\cdot g_{2,1}+\nu_1}\cdot\left(1+SINR_2\right)\right)\right) = 0 \tag{5.8}$$

$$\left(\left(\left(\frac{P_1\cdot g_{1,1}+P_2\cdot g_{2,1}+\nu_1}{P_2\cdot g_{2,1}+\nu_1}\right)\cdot\frac{-g_{1,2}\cdot P_2\cdot g_{2,2}}{\left(P_1\cdot g_{1,2}+\nu_2\right)^2}\right)+\left(\frac{g_{1,1}}{P_2\cdot g_{2,1}+\nu_1}\cdot\left(\frac{P_2\cdot g_{2,2}+P_1\cdot g_{1,2}+\nu_2}{P_1\cdot g_{1,2}+\nu_2}\right)\right)\right) = 0 \tag{5.9}$$

$$\left(\left(\frac{\left(P_1\cdot g_{1,1}+P_2\cdot g_{2,1}+\nu_1\right)\cdot\left(-g_{1,2}\cdot P_2\cdot g_{2,2}\right)}{\left(P_2\cdot g_{2,1}+\nu_1\right)\cdot\left(P_1\cdot g_{1,2}+\nu_2\right)^2}\right)+\left(\frac{g_{1,1}\cdot\left(P_2\cdot g_{2,2}+P_1\cdot g_{1,2}+\nu_2\right)}{\left(P_2\cdot g_{2,1}+\nu_1\right)\cdot\left(P_1\cdot g_{1,2}+\nu_2\right)}\right)\right) = 0 \tag{5.10}$$

$$\left(\left(\frac{\left(P_1\cdot g_{1,1}+P_2\cdot g_{2,1}+\nu_1\right)\cdot\left(-g_{1,2}\cdot P_2\cdot g_{2,2}\right)}{\left(P_2\cdot g_{2,1}+\nu_1\right)\cdot\left(P_1\cdot g_{1,2}+\nu_2\right)^2}\right)+\left(\frac{g_{1,1}\cdot\left(P_2\cdot g_{2,2}+P_1\cdot g_{1,2}+\nu_2\right)\cdot\left(P_1\cdot g_{1,2}+\nu_2\right)}{\left(P_2\cdot g_{2,1}+\nu_1\right)\cdot\left(P_1\cdot g_{1,2}+\nu_2\right)^2}\right)\right) = 0 \tag{5.11}$$

$$\left(\frac{\left(P_1\cdot g_{1,1}+P_2\cdot g_{2,1}+\nu_1\right)\cdot\left(-g_{1,2}\cdot P_2\cdot g_{2,2}\right)+g_{1,1}\cdot\left(P_2\cdot g_{2,2}+P_1\cdot g_{1,2}+\nu_2\right)\cdot\left(P_1\cdot g_{1,2}+\nu_2\right)}{\left(P_2\cdot g_{2,1}+\nu_1\right)\cdot\left(P_1\cdot g_{1,2}+\nu_2\right)^2}\right) = 0 \tag{5.12}$$

But since we know that $\left(P_2\cdot g_{2,1}+\nu_1\right)\cdot\left(P_1\cdot g_{1,2}+\nu_2\right)^2 > 0$ is always true, we can remove

the denominator since it does not affect the equality. Simplifying farther we get,

$$\begin{aligned}(P_1 \cdot g_{1,1} + P_2 \cdot g_{2,1} + \nu_1) \cdot (-g_{1,2} \cdot P_2 \cdot g_{2,2}) + \\ g_{1,1} \cdot (P_2 \cdot g_{2,2} + P_1 \cdot g_{1,2} + \nu_2) \cdot (P_1 \cdot g_{1,2} + \nu_2) = 0\end{aligned} \tag{5.13}$$

$$\begin{aligned}-P_1 \cdot g_{1,1} \cdot g_{1,2} \cdot P_2 \cdot g_{2,2} - P_2 \cdot g_{2,1} \cdot g_{1,2} \cdot P_2 \cdot g_{2,2} - \nu_1 \cdot g_{1,2} \cdot P_2 \cdot g_{2,2} + \\ g_{1,1} \cdot P_2 \cdot g_{2,2} \cdot P_1 \cdot g_{1,2} + g_{1,1} \cdot P_2 \cdot g_{2,2} \cdot \nu_2 + g_{1,1} \cdot P_1 \cdot g_{1,2} \cdot P_1 \cdot g_{1,2} + \\ g_{1,1} \cdot P_1 \cdot g_{1,2} \cdot \nu_2 + g_{1,1} \cdot \nu_2 \cdot P_1 \cdot g_{1,2} + g_{1,1} \cdot \nu_2 \cdot \nu_2 = 0\end{aligned} \tag{5.14}$$

$$\begin{aligned}-P_1 \cdot (g_{1,1} \cdot g_{1,2} \cdot P_2 \cdot g_{2,2}) - (P_2 \cdot g_{2,1} \cdot g_{1,2} \cdot P_2 \cdot g_{2,2}) - (\nu_1 \cdot g_{1,2} \cdot P_2 \cdot g_{2,2}) + P_1 \cdot \\ (g_{1,1} \cdot P_2 \cdot g_{2,2} \cdot g_{1,2}) + (g_{1,1} \cdot P_2 \cdot g_{2,2} \cdot \nu_2) + P_1^2 \cdot (g_{1,1} \cdot g_{1,2} \cdot g_{1,2}) + P_1 \cdot \\ (g_{1,1} \cdot g_{1,2} \cdot \nu_2) + P_1 \cdot (g_{1,1} \cdot \nu_2 \cdot g_{1,2}) + (g_{1,1} \cdot \nu_2 \cdot \nu_2) = 0\end{aligned} \tag{5.15}$$

$$\begin{aligned}P_1^2 \cdot (g_{1,1} \cdot g_{1,2} \cdot g_{1,2}) - P_1 \cdot (g_{1,1} \cdot g_{1,2} \cdot P_2 \cdot g_{2,2}) + P_1 \cdot (g_{1,1} \cdot P_2 \cdot g_{2,2} \cdot g_{1,2}) + \\ P_1 \cdot (g_{1,1} \cdot g_{1,2} \cdot \nu_2) + P_1 \cdot (g_{1,1} \cdot \nu_2 \cdot g_{1,2}) - (P_2 \cdot g_{2,1} \cdot g_{1,2} \cdot P_2 \cdot g_{2,2}) - \\ (\nu_1 \cdot g_{1,2} \cdot P_2 \cdot g_{2,2}) + (g_{1,1} \cdot P_2 \cdot g_{2,2} \cdot \nu_2) + (g_{1,1} \cdot \nu_2 \cdot \nu_2) = 0\end{aligned} \tag{5.16}$$

$$\begin{aligned}P_1^2 \cdot (g_{1,1} \cdot g_{1,2} \cdot g_{1,2}) + \\ P_1 \cdot (-g_{1,1} \cdot g_{1,2} \cdot P_2 \cdot g_{2,2} + g_{1,1} \cdot P_2 \cdot g_{2,2} \cdot g_{1,2} + g_{1,1} \cdot g_{1,2} \cdot \nu_2 + g_{1,1} \cdot \nu_2 \cdot g_{1,2}) + \\ (g_{1,1} \cdot P_2 \cdot g_{2,2} \cdot \nu_2 + g_{1,1} \cdot \nu_2 \cdot \nu_2 - P_2 \cdot g_{2,1} \cdot g_{1,2} \cdot P_2 \cdot g_{2,2} - \nu_1 \cdot g_{1,2} \cdot P_2 \cdot g_{2,2}) = 0\end{aligned} \tag{5.17}$$

$$\begin{aligned}P_1^2 \cdot (g_{1,1} \cdot g_{1,2} \cdot g_{1,2}) + P_1 \cdot (2 \cdot g_{1,1} \cdot g_{1,2} \cdot \nu_2) + \\ (g_{1,1} \cdot P_2 \cdot g_{2,2} \cdot \nu_2 + g_{1,1} \cdot \nu_2 \cdot \nu_2 - P_2 \cdot g_{2,1} \cdot g_{1,2} \cdot P_2 \cdot g_{2,2} - \nu_1 \cdot g_{1,2} \cdot P_2 \cdot g_{2,2}) = 0\end{aligned} \tag{5.18}$$

So to solve for P_1, if we define a, b and c as given in (5.19), then the equation above can be written as shown in (5.20) and the roots are given in (5.21).

$$\begin{aligned}a &= g_{1,1} \cdot g_{1,2} \cdot g_{1,2} \\ b &= 2 \cdot g_{1,1} \cdot g_{1,2} \cdot \nu_2 \\ c &= g_{1,1} \cdot P_2 \cdot g_{2,2} \cdot \nu_2 + g_{1,1} \cdot \nu_2 \cdot \nu_2 - P_2 \cdot g_{2,1} \cdot g_{1,2} \cdot P_2 \cdot g_{2,2} - \nu_1 \cdot g_{1,2} \cdot P_2 \cdot g_{2,2}\end{aligned} \tag{5.19}$$

$$P_1^2 \cdot a + P_1 \cdot b + c = 0 \tag{5.20}$$

$$roots = \frac{-b \pm \sqrt{b^2 - 4ac}}{2 \cdot a} \tag{5.21}$$

Because we know $a > 0$ and $b > 0$, we have at most a single positive root and this root is equal to $\frac{-b+\sqrt{b^2-4ac}}{2 \cdot a}$. To ensure this root is not a maximum, let us remember that if a root with respect to P_1 is a maximum then the slope is negative as P_1 increases past the root. If P_1 were to go to infinity with a constant P_2, $SINR_1$ would also go to infinity meaning our C_1 and C_{tot} would go to infinity also. This means that the slope is positive as P_1 increases past the possibly positive root. Since the slope is positive past this final root, this root cannot be a maximum since it does not match the definition of a maximum. Since no maximum values lie on the sum-rate curve with respect to P_1, the maximum sum-rate will occur either where $P_1 = P_{min}$ or $P_1 = P_{max}$. ■

5.4 Conclusions for Power Control

Our conclusions for the benefits of binary power control include:

- Binary (on/off) power control and full power control are often very close when optimizing sum-rate due to the fact that, in the optimal strategy with full power control, most users transmit at either 0 or P_{max}. We proved that for the two-user case it is in fact optimal for both users to transmit at either 0 or P_{max} when optimizing sum-rate (this case is also proven independently in [96]). These same authors show that the optimal result with FPC is generally equivalent to the optimal result with BPC for $N > 2$.
- Channel selection yields no benefit when there is only a single channel (as is expected). This benefit increases as the number of users increases and the benefit decreases as the number of channels increases past two channels.
- Binary power control has the most benefit over no power control when many users and few channels exist (specifically when not all users can be supported simultaneously without causing harmful interference to each other).
- The results suggest that full power control when optimizing sum of log-rate has a Jain fairness benefit (albeit possibly mild) over transmitting with no power control; however, there is a loss in sum-rate when using full power control versus no power control (we note that this is by design since we are optimizing fairness at the expense of sum-rate).
- In target-based SINR systems where we measure performance by the number of admitted users, using full power control has a benefit over binary power control. This is due to the fact that the benefit of collaboration results from minimizing interference. Since minimizing interference involves minimizing transmit power (while still achieving the target performance), the optimal transmit power will be less than P_{max} (as explained in [95]). For this reason, binary power control is unable to achieve the equal performance with full power control in multi-user systems with target-based optimization.

5.5 Value of Multiband Transmission

We now switch focus to the benefits of multi-band approaches. We specifically focus on the benefit that we can obtain through optimization by using spread spectrum or multi-channel approaches rather than the single-channel approach we have assumed up until now. For the case of $N = 6$, $M = 2$, $SNR_{min} = 3dB$, and $MUI = 5$, we plotted the sum-rate per user CDF for SC, MC, and SS in Figure 5.15. We notice immediately that spread spectrum

and single-channel approaches yield nearly equivalent performance, while multi-channel approaches tend to yield better performance. We note that the comparison of using spread spectrum techniques to using single-channel techniques is analogous to comparing the benefit of interference averaging over interference avoidance. However, spread spectrum does not tend to benefit performance in our scenario because the static channel we assume does not capture the dynamic nature of the channel. For example, in our model each user maximizes its own performance by transmitting full power on the channel that yields it the best performance; however, since our model does not consider that the channel gains may change, the optimization procedure does not consider that a more optimal result may spread the power among multiple channels (to average out the impact of channel gains decreasing). Multi-channel approaches can significantly increase throughput if the number of channels is increased (in some cases doubling the number of channels can nearly double its capacity).

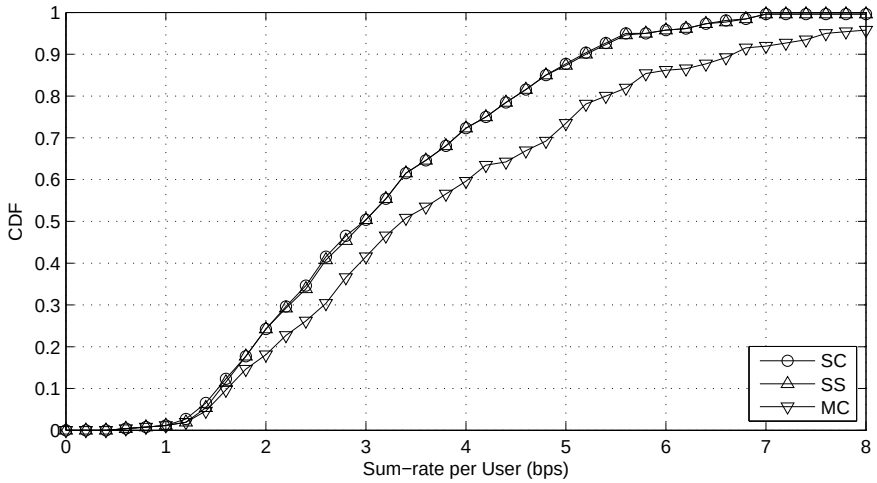

Figure 5.15: CDF of Average Sum-rate per User for Multi-band Approaches

For the case of $SNR_{min} = 3dB$ and $MUI = 5$, we plot the ratio of the average sum-rate of SS techniques divided by the average sum-rate of SC techniques against the number of users and channels in Figure 5.16. There is practically no benefit for using spread spectrum techniques over single-channel techniques for our static system model. For the same case, we plot the ratio of the average sum-rate of MC techniques divided by the average sum-rate of SC techniques against the number of users and channels in Figure 5.17. Obviously, as we increase the number of channels, there is a larger benefit to using MC techniques over SC techniques. Furthermore, the benefit to using MC techniques increases even further when the number of users is less than the number of channels. This is because the MC approach is able to utilize all the bandwidth by giving users multiple channels while the SC technique is only able to give each user a single channel.

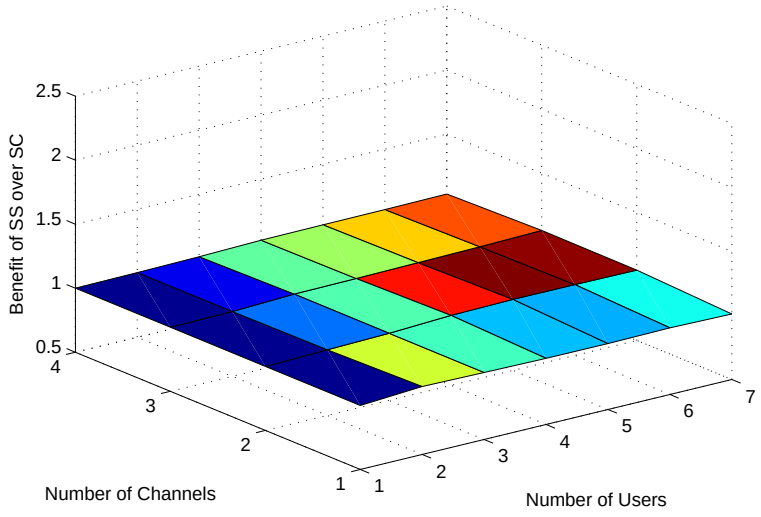

Figure 5.16: Benefit of Using Spread Spectrum (SS) over Single-Channel (SC)

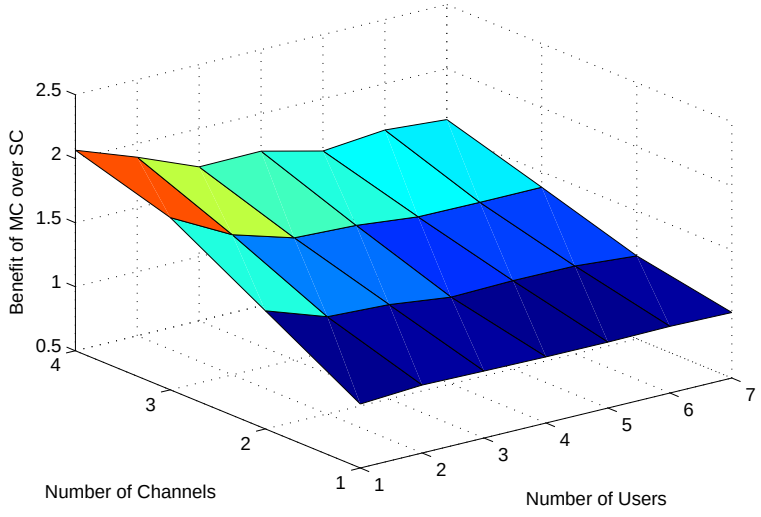

Figure 5.17: Benefit of Using Multi-channel (MC) over Single-Channel (SC)

5.6 Proof for Spread Spectrum Two-User Case

In this section, we explore the equivalence of optimal spread spectrum and single-channel techniques in the specific case of optimizing sum-rate with two users and two channels. We make few assumptions about noise powers (for each channel and user), channel gains (from a transmitter to a receiver on a specific channel), and transmit powers (for each user) to maintain generality in this analysis. We derive the general sum-rate expression in Section 5.6.1. In Section 5.6.2, we prove that if the optimal result involves one user using a single-channel technique, then both users must use single-channel techniques to achieve the optimal result. This means that the optimal result involves either both users using spread spectrum or neither user using spread spectrum. In Section 5.6.3, we continue this analysis to consider specific situations in which the optimal spread spectrum and single-channel results are equivalent. Finally, we consider one more set of cases for which we cannot prove that optimal spread spectrum and single-channel results are equivalent, but through simulation results, we show that we are unable to find a set of input variables in which optimal spread spectrum techniques outperform optimal single channel techniques.

5.6.1 Sum-rate Derivation

We maximize sum-rate by optimizing over the channel selection fractions (i.e., the fraction of power transmitted on each channel). The channel selection fraction ρ_i tells us what fraction of user i's power, P_i, is transmitted on channel 1. The noise term, $\nu_{i,k}$, is the noise power seen at receiver i on channel k (we note that we allow the noise powers to be different for each user on each channel to account for the possibility of outside Gaussian interference). The channel gain, $g_{i,j,k}$, is the gain from transmitter i to receiver j on channel k. The sum-rate is calculated as shown in (5.22).

$$\begin{aligned} C_{tot} &= C_1 + C_2 && (5.22)\\ &= log_2\left(1 + SINR_1\right) + log_2\left(1 + SINR_2\right) && (5.23)\\ &= log_2\left(\left(1 + SINR_1\right) \cdot \left(1 + SINR_2\right)\right) && (5.24) \end{aligned}$$

Since we maximize sum-rate, this is the same as maximizing,

$$\left(1 + SINR_1\right) \cdot \left(1 + SINR_2\right) \tag{5.25}$$

Using the transmit powers, channel gains, and noise powers and assuming that we use maximum ratio combining, we can represent the optimization equation as,

$$\left(1 + \frac{P_1 \cdot \rho_1 \cdot g_{1,1,1}}{P_2 \cdot \rho_2 \cdot g_{2,1,1} + \nu_{1,1}} + \frac{P_1 \cdot (1-\rho_1) \cdot g_{1,1,2}}{P_2 \cdot (1-\rho_2) \cdot g_{2,1,2} + \nu_{1,2}}\right) \cdot \left(1 + \frac{P_2 \cdot \rho_2 \cdot g_{2,2,1}}{P_1 \cdot \rho_1 \cdot g_{1,2,1} + \nu_{2,1}} + \frac{P_2 \cdot (1-\rho_2) \cdot g_{2,2,2}}{P_1 \cdot (1-\rho_1) \cdot g_{1,2,2} + \nu_{2,2}}\right)$$

5.6.2 Both or Neither Spread Spectrum

In this section, we prove by contradiction that to achieve optimal sum-rate either both users or neither user uses spread spectrum techniques. Without loss of generality, we will assume that achieving optimal sum-rate requires user 1 to transmit full power on channel 1. If we can prove that optimal sum-rate results only when $\rho_2 = 0$ or $\rho_2 = 1$, then we have proven that if the optimal result involves one user transmitting on a single channel, the optimal result requires that both users transmit on a single channel.

To prove this, we take the derivative of the sum-rate with respect to ρ_2 in (5.26). If we can prove that no maximum exists as $0 < \rho_2 < 1$, then the optimal sum-rate will be either where $\rho_2 = 0$ or $\rho_2 = 1$. The derivative is simplified in 5.27 and 5.28.

$$\frac{\delta C_{tot}}{\delta \rho_2} = (1 + SINR_1) \cdot \frac{\delta SINR_2}{\delta \rho_2} + \frac{\delta SINR_1}{\delta \rho_2} \cdot (1 + SINR_2) \tag{5.26}$$

$$= \frac{\rho_2^2 \left(g_{2,1,1}^2 \cdot g_{2,2,1} \left(\nu_{2,2} - I\right)\right) + \rho_2 \left(2 \cdot \nu_{1,1} g_{2,1,1} g_{2,2,1} \left(\nu_{2,2} - I\right)\right)}{\left(\rho_2 \left(g_{2,1,1} \nu_{2,2} \left(P_1 g_{1,2,1} + \nu_{2,1}\right)\right) + \left(\nu_{1,1} \nu_{2,2} \left(P_1 g_{1,2,1} + \nu_{2,1}\right)\right)\right)^2} + \tag{5.27}$$

$$\frac{\left(\nu_{1,1} C g_{2,2,1} \left(\nu_{2,2} - I\right) - P_1 g_{1,1,1} g_{2,1,1} I \left(\nu_{2,2} + g_{2,2,1}\right)\right)}{\left(\rho_2 \left(g_{2,1,1} \nu_{2,2} \left(P_1 g_{1,2,1} + \nu_{2,1}\right)\right) + \left(\nu_{1,1} \nu_{2,2} \left(P_1 g_{1,2,1} + \nu_{2,1}\right)\right)\right)^2} \tag{5.28}$$

where

$$I = P_1 g_{1,2,1} + \nu_{2,1} \tag{5.29}$$

$$C = P_1 g_{1,1,1} + \nu_{1,1} \tag{5.30}$$

We will then solve for when this derivative is equal to zero since this will tell us where maximums may occur. We note that the denominator is always positive, so we only need to solve for when the numerator is equal to zero as shown in (5.31). This results in three cases of interest, namely $\nu_{2,2} < I$, $\nu_{2,2} = I$ and $\nu_{2,2} > I$. If we define a, b, and c as given in (5.32), our roots are defined as shown in (5.33).

$$\begin{aligned} &\rho_2^2 \left(g_{2,1,1}^2 \cdot g_{2,2,1} \left(\nu_{2,2} - I\right)\right) + \rho_2 \left(2 \nu_{1,1} g_{2,1,1} g_{2,2,1} \left(\nu_{2,2} - I\right)\right) + \\ &\quad \left(\nu_{1,1} C g_{2,2,1} \left(\nu_{2,2} - I\right) - P_1 g_{1,1,1} g_{2,1,1} I \left(\nu_{2,2} + g_{2,2,1}\right)\right) = 0 \end{aligned} \tag{5.31}$$

$$\begin{aligned} a &= g_{2,1,1}^2 \cdot g_{2,2,1} \left(\nu_{2,2} - I\right) \\ b &= 2 \nu_{1,1} g_{2,1,1} g_{2,2,1} \left(\nu_{2,2} - I\right) \\ c &= \nu_{1,1} C g_{2,2,1} \left(\nu_{2,2} - I\right) - P_1 g_{1,1,1} g_{2,1,1} I \left(\nu_{2,2} + g_{2,2,1}\right) \end{aligned} \tag{5.32}$$

$$roots = \frac{-b \pm \sqrt{b^2 - 4 \cdot a \cdot c}}{2 \cdot a} \tag{5.33}$$

If $\nu_{2,2} < I$, then $a < 0$ $b < 0$ and $c < 0$. This means that $\sqrt{b^2 - 4 \cdot a \cdot c} < |b|$. This means that the numerator of (5.33) must be positive and the denominator must be negative; therefore all roots must be negative, so no zeros exist in the range $0 < \rho_2 < 1$. If no zeros exist in this range, then no maximums may exist in this range and consequently the optimal choice will lie either at $\rho_2 = 0$ or $\rho_2 = 1$.

If $\nu_{2,2} = I$, then $a = 0$, $b = 0$, and $c \leq 0$. If $c = 0$, then the slope is equal to zero for all values of ρ_2. We note that conceptually, $\nu_{2,2} = I$ means that the equivalent Gaussian noise plus interference on both channels is equal. Accordingly, the user may select either channel and achieve equivalent performance for himself. However, the derivative is typically negative because as user 2 places more power onto channel 1, he increases the interference seen by user 1 and thus hurts sum-rate. We note that in the equations, the case where the derivative is equal to zero is when $P_1 = 0$ (User 1 transmits no power and thus has no capacity), $g_{1,1,1} = 0$ (even if user 1 transmits power on channel one, no power reaches the receiver), or $g_{2,1,1} = 0$ (the case in which user 2 does not interfere with user 1 at all).

Thus when the derivative is negative (due to user 2 causing harmful interference to user 1), the optimal performance for user 2 is when $\rho_2 = 0$ or user 2 places all its power on channel 2. When the derivative is zero everywhere, user 2 may place any fraction of his power on either channel, but optimal performance does not require it to use spread spectrum. Optimal performance still occurs at $\rho_2 = 0$ or $\rho_2 = 1$, but it can also exist for $0 < \rho_2 < 1$.

If $\nu_{2,2} > I$, then $a > 0$ and $b > 0$. This means that only one root may be positive (as was shown previously in Section 5.3). If this positive root is a maximum, then we know $\frac{\delta C_{tot}}{\delta \rho_2} < 0$ for $root < \rho_2 < \infty$. On the contrary though, if ρ_2 goes toward infinity (we note that this concept is based on theoretical math and not a real world scenario as $\rho_2 \leq 1$ in all real scenarios), $\frac{\delta C_{tot}}{\delta \rho_2}$ goes to a positive constant (as given in (5.34)). Therefore no maximum may exist in the range $0 < \rho_2 < 1$ so the optimal choice for ρ_2 will be either $\rho_2 = 0$ or $\rho_2 = 1$.

$$\lim_{\rho_2 \to \infty} \frac{\delta C_{tot}}{\delta \rho_2} = \frac{\left(g_{2,1,1}^2 \cdot g_{2,2,1} \left(\nu_{2,2} - I\right)\right)}{\left(g_{2,1,1} \nu_{2,2} \left(P_1 g_{1,2,1} + \nu_{2,1}\right)\right)^2} \tag{5.34}$$

Therefore for all three cases, we have shown that the optimal performance always occurs for either $\rho_2 = 0$ or $\rho_2 = 1$. For the case of $\nu_{2,2} = I$, we can use spread spectrum to obtain optimal sum-rate (in the special case where user 2 does not degrade the performance of user 1 by transmitting on channel 1), but no benefit exists for using spread spectrum techniques over single-channel techniques since single-channel techniques can also achieve the optimal sum-rate.

Since the second user achieves optimal performance when using single-channel techniques, we have proven that if one user uses single-channel techniques that it is optimal for both users to use single-channel techniques. This means that either both users use spread spectrum techniques or neither user uses spread spectrum techniques.

5.6.3 Cases Where Spread Spectrum is Not Used

Users Select Channel with Best SINR

For spread spectrum techniques, optimizing the capacity of a user is equivalent with optimizing its $SINR$. The $SINR$ for a spread spectrum user transmitting on two channels is given by 5.35. (A derivation of the total $SINR$ using a spread spectrum technique can be found in Appendix A).

$$\begin{aligned} SINR &= \frac{P_1 \cdot g_{1,1,1} \cdot \rho}{\nu_{1,1}} + \frac{P_1 \cdot g_{1,1,2} \cdot (1-\rho)}{\nu_{1,2}} && (5.35) \\ &= P_1 \left(\rho \cdot \frac{g_{1,1,1}}{\nu_{1,1}} + (1-\rho) \cdot \frac{g_{1,1,2}}{\nu_{1,2}} \right) && (5.36) \\ &= P_1 \left(\frac{g_{1,1,2}}{\nu_{1,2}} + \rho \cdot \frac{g_{1,1,1}}{\nu_{1,1}} - \rho \cdot \frac{g_{1,1,2}}{\nu_{1,2}} \right) && (5.37) \\ &= P_1 \left(\frac{g_{1,1,2}}{\nu_{1,2}} + \rho \cdot \left(\frac{g_{1,1,1}}{\nu_{1,1}} - \frac{g_{1,1,2}}{\nu_{1,2}} \right) \right) && (5.38) \end{aligned}$$

To determine what value of ρ will maximize the $SINR$ we need to know whether $\frac{g_{1,1,2}}{\nu_{1,2}}$ or $\frac{g_{1,1,1}}{\nu_{1,1}}$ is larger. If $\frac{g_{1,1,2}}{\nu_{1,2}} > \frac{g_{1,1,1}}{\nu_{1,1}}$ then decreasing ρ increases the $SINR$ making the optimal $\rho = 0$ (we call this preferring channel 2). If $\frac{g_{1,1,2}}{\nu_{1,2}} < \frac{g_{1,1,1}}{\nu_{1,1}}$ then increasing ρ increases the $SINR$ making the optimal $\rho = 1$ (we call this preferring channel 1). If $\frac{g_{1,1,2}}{\nu_{1,2}} = \frac{g_{1,1,1}}{\nu_{1,1}}$ then changing ρ has no impact on the $SINR$ meaning that there is no benefit to spread spectrum techniques (we call this preferring neither channel). Conceptually, this means that if a user can obtain a better $SINR$ on one channel, the user will achieve the best individual performance by placing all power on that channel. However, if the user obtains the same $SINR$ on both channels, the user can put any fraction of its power on either channel and obtain the optimal performance (and ρ has no impact on individual performance).

We note that if a second user were present on the same two channels this only affects the optimal performance of the user being considered by affecting the $\nu_{1,1}$ and $\nu_{1,2}$ terms. For example, it is possible that the second user changing its own channel selection fraction, ρ_2, could cause the original user to favor a different channel.

Separation of Cases

In this section, we study specific situations and determine whether a benefit exists for spread spectrum techniques. We put each user into one of three groups based upon channel preference (i.e., preferring channel 1, preferring channel 2, or preferring neither channel). Between the two users, there exist nine possible states. The optimal SC channel selections for each

Table 5.1: Optimal Channel Selection Assuming Single-Channel. Given the channel preferences of users 1 and 2, the (i,j) pairs represent the possible optimal channel selection solutions where i is the channel user 1 selects and j is the channel user 2 selects.

	User 1, Channel 1	User 1, No Preference	User 1, Channel 2
User 2, Channel 1	(1,1) (1,2) (2,1)	(2,1)	(2,1)
User 2, No Preference	(1,2)	(1,2) (2,1)	(2,1)
User 2, Channel 2	(1,2)	(1,2)	(1,2) (2,1) (2,2)

of these states (assuming users interfere with each other) are shown in Table 5.1.

We will first prove that single-channel techniques achieve optimal performance for the case in which one user prefers a channel and the other user does not prefer that same channel (the other user may prefer the other channel or neither channel). We will second prove that single-channel techniques achieve optimal performance for the case in which neither user prefers a channel. Between this set of proofs, the only remaining case is that in which both users prefer the same channel. We explore this case in the final subsection through simulation.

Users Do Not Prefer the Same Channel

If both users do not prefer the same channel, we prove here that the optimal performance occurs using single-channel techniques. We first note that the optimal sum-rate cannot be more than the summation of the optimal capacity of the two individual users. With this being the case, we can say that if both users prefer separate channels and both receive the channel that they prefer that the optimal total performance is achieved since the optimal individual performances are achieved. Similarly, if one user prefers a channel and the other user has no preference, the user who prefers a channel can optimize his own capacity by using the channel he prefers and the second user can optimize his performance by using the other channel. Therefore, if one user prefers a channel, and the other user does not prefer that same channel, optimal capacity is achieved by single-channel techniques.

No Preference for Either Channel

Similar to the previous section, the optimal total capacity cannot be greater than the sum of the optimal performance of the individual users. Since we can achieve the optimal performance of both individual users by placing users on separate channels, again we have no benefit of spread spectrum techniques.

Simulation Result for Remaining Cases

We begin by noting that the results from the previous section have shown that the optimal sum-rate of spread spectrum has no benefit over the optimal sum-rate of single-channel when a user prefers a channel and the other user does not prefer the same channel and when neither user prefers a channel. The only case that still has to be proven is the case in which both users prefer the same channel. While this case is left unproven analytically, through simulation we show that we are unable to find any circumstance in which single-channel techniques are incapable of achieving the optimal performance.

We simulated 1000 cases in which both users preferred the same channel (for our simulations we made this channel 1) and determined the optimal channel selection fraction (ρ) for each user. We plotted a PDF of the channel selection fractions in Figure 5.18. We notice that, because the channel selection fraction is always 0 or 1, in no simulations was the optimal result achieved by spread spectrum techniques. While this result does not guarantee that spread spectrum techniques cannot be more optimal than single-channel techniques, it does show that in the majority of situations that we consider we do not expect the performance of spread spectrum techniques to surpass the performance of single-channel techniques. We can further see that when both users prefer channel 1 that about 15% of the time both users transmit on channel 1 while about 85% of the time users transmit on separate channels.

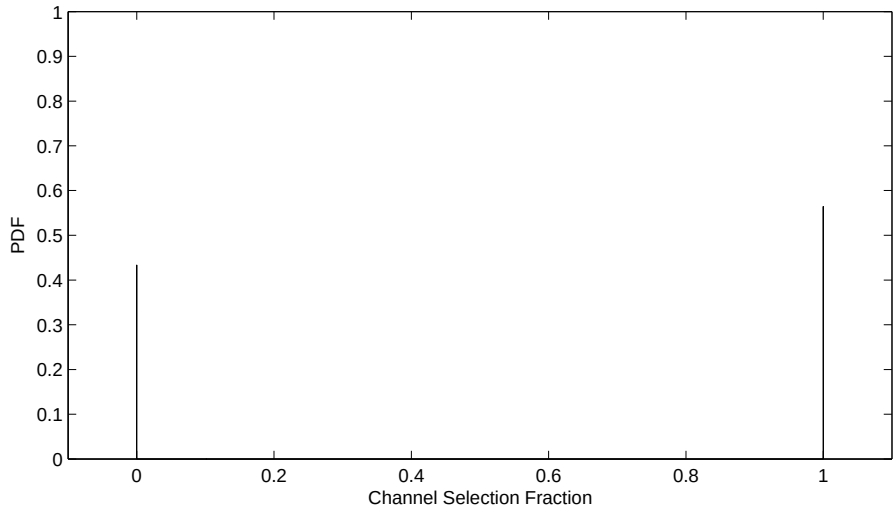

Figure 5.18: PDF of Channel Selection Fraction ρ over 1000 Simulations

5.7 Multi-channel Example Case

In this section, we explore how the benefit of using multi-channel techniques over single-channel techniques are related to the number of channels and users. We explore the reasoning through example situations where we show the reasoning behind the gains. We first will present three examples and simply state the results of the optimization using single-channel and multi-channel. We will then conclude about the reasoning behind what impact the number of users and the number of channels have on the benefit of multi-channel approaches over single-channel approaches.

We note that because our noise value in these sample cases of 0.5 is unrealistic to real life scenarios we must scale the receive signal power accordingly to achieve the SNR_{min}. This means that while we may have gains that are larger than 1 these gains do not represent a gain in power from the Tx node to the Rx node. They simply represent a relative gain to scale the Rx power to the same scale as the noise power.

5.7.1 Example 1

For all sample cases, we will assume 5 users and 2 channels. We also assume a gain matrix that is created only using shadowing that is perfectly correlated across channels. Generated by our simulation program, the channel gain matrix **g** for example 1 is given by,

$$g = \begin{bmatrix} 0.2568 & 0.1732 & 0.0959 & 7.6531 & 0.0072 \\ 27.8531 & 17.1953 & 0.0609 & 0.0073 & 5.8597 \\ 0.0241 & 0.1132 & 11.968 & 0.1905 & 307.5867 \\ 0.1366 & 0.0254 & 0.569 & 0.0273 & 0.7363 \\ 2.1831 & 0.4667 & 10.2892 & 0.0641 & 0.3911 \end{bmatrix} \tag{5.39}$$

We note immediately that users 2 and 3 have the highest gains and thus have the best channels. The optimal single-channel results from user 1 transmitting full power on channel 2, user 2 transmitting full power on channel 1, user 3 transmitting full power on channel 1, user 4 not transmitting and user 5 transmitting full power on channel 2. The optimal multi-channel results from users 1, 4, and 5 not transmitting, and users 2 and 3 splitting their power evenly between the two channels. This makes the optimal capacity vector for single-channel,

$$C_{SC} = \begin{bmatrix} 0.0859 & 3.1917 & 2.767 & 0 & 0.2568 \end{bmatrix} \tag{5.40}$$

and the optimal capacity vector for multi-channel,

$$C_{MC} = \begin{bmatrix} 0 & 4.7462 & 3.9616 & 0 & 0 \end{bmatrix} \tag{5.41}$$

Therefore in this case, the total sum-rate increases from 6.3015 to 8.7078.

5.7.2 Example 2

For example 2, we use the gain matrix,

$$g = \begin{bmatrix} 0.8974 & 88.1618 & 203.9627 & 7.8183 & 3556.8 \\ 0.6786 & 0.2873 & 0.1315 & 0.7699 & 32.9 \\ 33.7878 & 3.5097 & 74.0855 & 0.0252 & 69.2 \\ 17.0856 & 0.0673 & 0.0686 & 0.1363 & 0.4 \\ 1.9581 & 0.0467 & 0.0003 & 0.1841 & 2.5 \end{bmatrix} \tag{5.42}$$

We note that for this example, the dominant user is user 3 with a channel gain of around 74. User 5 is the next most dominant user with a channel gain of 2.5. In this case, the optimal single-channel result causes users 1 and 2 to not transmit, users 3 and 4 to be on channel 2 and user 5 to be on channel 1. The optimal multi-channel result causes users 1 and 2 to not transmit and users 3, 4, and 5 to transmit evenly on both channels. This makes the optimal capacity vector for single-channel,

$$C_{SC} = \begin{bmatrix} 0 & 0 & 5.2022 & 0.094 & 1.1721 \end{bmatrix} \tag{5.43}$$

and the optimal capacity vector for multi-channel,

$$C_{MC} = \begin{bmatrix} 0 & 0 & 8.5272 & 0.0919 & 0.0966 \end{bmatrix} \tag{5.44}$$

Therefore in this case, the total sum-rate increases from 6.4683 to 8.7158.

5.7.3 Example 3

For example 3, we use the gain matrix,

$$g = \begin{bmatrix} 23.2064 & 0.0242 & 1.1104 & 104.7119 & 0.1302 \\ 0.3255 & 4.1085 & 0.8065 & 0.1326 & 0.0698 \\ 3.5644 & 3.8766 & 97.0498 & 0.0302 & 2.453101 \\ 5.1708 & 0.1035 & 135.9654 & 0.02 & 0.894 \\ 2.5259 & 0.3343 & 1.161 & 1.5048 & 17.6882 \end{bmatrix} \tag{5.45}$$

In this example, user 3 has the best channel and users 1 and 5 both having similar average channels. The optimal single-channel sum-rate results from users 1, 2, and 5 transmitting on channel 1, user 3 transmitting on channel 2 and user 4 not transmitting. The optimal multi-channel results from users 1, 2, and 3 both transmitting evenly on channels 1 and 2 and users 4 and 5 not transmitting. This makes the optimal capacity vector for single-channel,

$$C_{SC} = \begin{bmatrix} 2.5319 & 1.4552 & 5.6301 & 0 & 3.1763 \end{bmatrix} \tag{5.46}$$

and the optimal capacity vector for multi-channel,

$$C_{MC} = \begin{bmatrix} 3.9573 & 1.2081 & 8.2424 & 0 & 0 \end{bmatrix} \tag{5.47}$$

Therefore in this case, the total sum-rate increases from 12.7936 to 13.4078.

5.7.4 Findings

We note first that in all three examples the user with the best channel always transmits. Furthermore, in all three examples some users with the worst channels were not allowed to transmit on either channel. This is related to the binary power control we discussed previously (the users with the best channels are given the channels they desire and other users are only allowed to use the same channels if they do not cause harmful interference to the users with better channels). Since both single-channel and multi-channel techniques tend to favor the transmissions of the users with the best channels, the optimal sum-rate of both tend to scale equivalently as the number of users increases.

On the contrary, additional channels causes the sum-rate of the multi-channel approach to increase more than the sum-rate of the single-channel approach for all three examples. This results from the users with the best channel gains being given the ability to use multi-channel approaches to better utilize the channels that were previously used by users that did not have the best channel gains. This can be seen in Example 1 by the fact that the capacities of users 2 and 3 (the two users with the best channels) increase when multi-channel is used, whereas users 1 and 5 are no longer allowed to transmit. In Example 2, the capacity of user 3 (the user with the best channel gain) increases when multi-channel is used while the capacity of users 4 and 5 decreases. In Example 3, the capacities of users 1 and 3 (again the users with the best channels) while the capacities of users 2 and 5 decrease. These show that the increase in capacity comes from the user with the best channel taking advantage of more channels and increasing the utility of the channel beyond what the users with worse channels were able to accomplish. This further shows that optimizing sum-rate with multi-channel techniques is expected to yield worse fairness than optimizing sum-rate with single-channel techniques.

5.7.5 Gains from Using Multi-channel Approaches

We note that the gains resulting from using multiple channels are related to the gains that individual users achieve for using multiple channels. A given user i can achieve the capacity of,

$$C_i = log_2 \left(1 + SNR_i\right) \tag{5.48}$$

Assuming this user can achieve the same SNR on all channels, the capacity this user can achieve by using multi-channel approaches with M channels is,

$$C_i = M \cdot log_2 \left(1 + \frac{SNR_i}{M}\right) \tag{5.49}$$

We note that the derivative of capacity with respect to SNR is,

$$\frac{\partial C_i}{\partial SNR_i} = M \cdot \frac{1}{1+\frac{SNR_i}{M}} \cdot \frac{1}{M} = \frac{1}{1+\frac{SNR_i}{M}} \tag{5.50}$$

So for any $SNR_i > 0$, as M increases, so does the derivative, $\frac{\partial C_i}{\partial SNR_i}$. Since all $C_i = 0$ when $SNR_i = 0$ regardless of the value of M, we can also say that, because the derivative of capacity with respect to SNR is larger for larger values of M, using more channels will always increase capacity. Furthermore, we can say that a larger SNR_i causes multi-channel approaches to yield a larger capacity increase (assuming $M > 1$). To determine the impact of the number of channels on performance benefit, we take the derivative of capacity with respect to M.

$$\frac{\partial C_i}{\partial M} = log_2\left(1+\frac{SNR_i}{M}\right) - ln(2) \cdot \frac{SNR_i}{M+SNR_i} \tag{5.51}$$

$$\lim_{M\to\infty} \frac{\partial C_i}{\partial M} = 0 \tag{5.52}$$

Taking the double derivative of capacity with respect to M we get,

$$\begin{aligned}
\frac{\partial^2 C_i}{\partial M^2} &= ln(2) \cdot \frac{SNR_i}{(M+SNR_i)^2} - ln(2) \cdot \frac{SNR_i}{M \cdot (M+SNR_i)} && (5.53)\\
&= ln(2) \cdot SNR_i \cdot \left(\frac{1}{(M+SNR_i)^2} - \frac{1}{M \cdot (M+SNR_i)}\right) && (5.54)\\
&= ln(2) \cdot SNR_i \cdot \left(\frac{1}{M^2+2 \cdot M \cdot SNR_i + SNR_i^2} - \frac{1}{M^2+M \cdot SNR_i}\right) && (5.55)\\
& && (5.56)
\end{aligned}$$

We note that,

$$\begin{aligned}
M \cdot SNR_i + SNR_i^2 &> 0 && (5.57)\\
M^2+2 \cdot M \cdot SNR_i + SNR_i^2 &> M^2+M \cdot SNR_i && (5.58)\\
\frac{1}{M^2+2 \cdot M \cdot SNR_i + SNR_i^2} &< \frac{1}{M^2+M \cdot SNR_i} && (5.59)\\
\frac{1}{M^2+2 \cdot M \cdot SNR_i + SNR_i^2} - \frac{1}{M^2+M \cdot SNR_i} &< 0 && (5.60)\\
\frac{\partial^2 C_i}{\partial M^2} &< 0 && (5.61)
\end{aligned}$$

If we combine the fact that the double derivative of capacity with respect to M is negative (assuming $M > 0$) and the derivative of capacity with respect to M as M approaches infinity is zero, we know that the derivative of capacity with respect to M is positive for $M > 0$. The fact that the derivative of capacity with respect to M is positive reinforces our comment that using more channels will always increase capacity. The fact that the double derivative of

capacity with respect to M is negative tells us that we should expect the benefit of additional channels to follow the law of diminishing returns.

The ratio of the capacity using multi-channel approaches with 2, 3 and 4 channels to the capacity using a single-channel approach is plotted in figure 5.19 for different SNR values. In this figure, we can see that larger values of M cause the benefit of using multiple channels to be larger by the fact that the curve with 4 channels yields a higher benefit than the curve with 3 channels (similarly the curve with 3 channels has more benefit than the curve with 2 channels). Also, as SNR_i increases, the benefit of using multi-channel approaches increases. We can also see that, as we add further channels, the benefit of using multiple channels follows the law of diminishing returns (a user with an $SNR = 10dB$ has a 50% gain from the second channel, but the third channel only adds and additional 35% and the fourth channel adds an additional 25%).

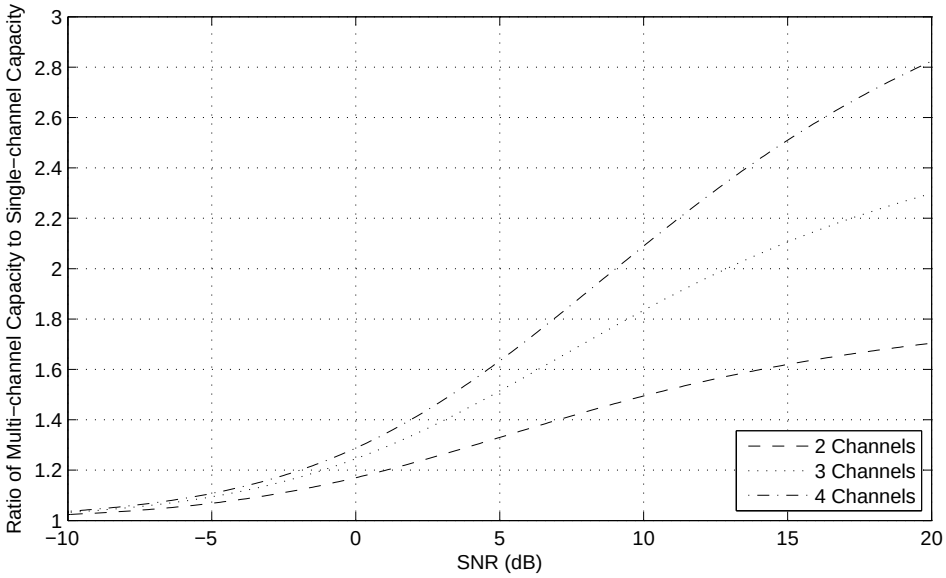

Figure 5.19: Benefit of MC over SC for a Single User

5.8 Conclusions for Multiband Approaches

We make the following conclusions for multi-band approaches:

- There is practically no benefit to using spread spectrum techniques compared to single-channel techniques using our system model. This is partially due to the fact that the optimization routine is given static channel gains and does not consider that the channel may be dynamic.
- Specifically for the case in which two users share two channels, we showed that for most scenarios the optimal result involves single-channel approaches.
- The benefit of multi-channel techniques compared to single-channel techniques tends to increase based on the number of channels present, but this follows the law of diminishing returns.
- This benefit seems to be affected very little if at all by the number of users collaborating. This is likely due to the fact that additional users only improves the multi-user diversity which affects SC and MC approaches equivalently (and thus the benefit of MC techniques over SC techniques does not change). From the multi-channel analysis, we can see that opening parallel data pipes is beneficial as compared to putting more power in a single data pipe. This is not surprising since this result is well known for MIMO systems.

Chapter 6

Distributed Techniques to Approach the Optimal

In the previous three chapters, we studied the impact of different parameters on the performance of collaborative and non-collaborative techniques. While these analyses assumed that sensing was our non-collaborative technique, we now consider whether alternative non-collaborative or minimally collaborative approaches might be able to achieve better performance than sensing. We begin the chapter by studying how the collaborative techniques achieve better performance than the non-collaborative technique of sensing. From these insights, we design two minimally collaborative techniques which attempt to eliminate the factors that cause the performance of sensing to be sub-optimal (e.g., sensing has no power control), and we show the performance for each technique. We then determine the level of collaboration that each of these techniques require, if any, and conclude by studying the impact of bursty transmission on each technique.

6.1 Interference Symmetry

In this section, we describe how a concept which we will call *interference symmetry* relates to the benefit of collaboration. Interference symmetry is the measure of how accurately a user a can predict the interference it would cause to user b by measuring the interference user b causes to user a. The most common example of interference symmetry is the hidden node and exposed node issues that arose with the introduction of random access protocols [97]. In this example, performance suffers either when a node assumes that it will interfere with another node but does not or does not assume it would interfere with another node but does. Throughout this section, we use the term direct channel gains for values of $g_{i,j,k}$ where $i = j$ and interference channel gains for the values of $g_{i,j,k}$ where $i \neq j$. Conceptually this

means that direct channel gains are the gains from a Tx node to its intended Rx node, and interference channel gains are the gains from a Tx node to all other Rx nodes.

We begin the section by introducing a metric which we use to place a value on the symmetry of the interference channel gains. We then examine the scenario of a centralized base station topology where the interference is symmetric and show that techniques exists that can achieve the optimal sum-rate in a distributed way. We then introduce some sample cases that illustrate why asymmetry hurts performance, and we conclude by plotting a correlation between the interference symmetry metric and the collaboration gain.

6.1.1 Interference Symmetry Metric

The term symmetry has been used in many similar (but not identical) ways in the literature (i.e., [98], [77], and [84]). In [98] and [99], symmetry is defined as the case when all cross channel gains are equal and all direct channel gains are equal. In [77], asymmetry is said to result from different cross channel gains or a disparity in transmission powers. In [84], symmetry is defined as the cross channel gains and direct channel gains from a Tx node being equal (the direct and interference channel gains are all equal for any particular transmitter to the receiver).

To be able to determine the benefit of collaboration for different levels of symmetry, we need to define a symmetry metric in order to compare different scenarios. We define our symmetry metric for a system as

$$M_{sym} = \frac{\sum_{i=1}^{N} \sum_{j=1,j\neq i}^{N} min\left(g_{i,j} \cdot g_{j,j}, g_{i,i} \cdot g_{j,i}\right)}{\sum_{i=1}^{N} \sum_{j=1,j\neq i}^{N} max\left(g_{i,j} \cdot g_{j,j}, g_{i,i} \cdot g_{j,i}\right)} \tag{6.1}$$

We note that, because $g_{i,j} > 0$ for all i and j, we know that $0 \leq M_{sym} \leq 1$. A channel is said to be 'perfectly symmetric' if $M_{sym} = 1$. For each pair of nodes in the system, this metric compares the value of $g_{i,j} \cdot g_{j,j}$ with $g_{i,i} \cdot g_{j,i}$. As these two values become further apart, the metric moves closer to zero, and as the two values become closer, the metric moves closer to one. We notice also that the larger the gain terms are for a given pair, the more the symmetry of that pair affects the system symmetry metric (and conversely smaller gain terms cause the symmetry of the pair to affect the system symmetry metric less). We will explain in more detail below how this symmetry impacts the value of collaboration.

We note that our definition of a perfectly symmetric channel holds in [98] and [84] but only holds for some cases in [77]. This is because [77] considers the asymmetry of maximum transmit powers (which we neglect because we assume a common maximum transmit power). However, our symmetry metric is consistent with that in [77] when neglecting transmit power asymmetry.

6.1.2 Multiaccess Channel Asymmetry

One of the most common scenarios that is perfectly symmetric according to our symmetry metric is the uplink of a network with a central base station. The optimal capacity of this scenario is shown in [84] to be achieved by using successive addition and cancellation of individual transmitters where the individual transmitters use water-filling (note we say optimal capacity instead of optimal sum-rate because this work optimizes the information theoretic capacity; information theoretic capacity allows receivers to use advanced techniques to mitigate or remove multi-user interference whereas our sum-rate metric assumes that all interference power is simply treated as noise). It is noted in [84] that the optimal capacity algorithm is inherently greedy in that the addition of each user involves that user using water-filling to achieve the best personal gain. Early works introduced an approach to achieve near optimal capacity using minimal feedback in [100] using just a few bits and [101] using just one bit. It is noted in [102] that maximizing total capacity in a CDMA cellular uplink is equivalent to minimizing the total squared correlation. This work introduces a fully distributed algorithm (no collaboration) to minimize total squared correlation and thus maximize total capacity, and in [103] the algorithm is proven to converge to the globally optimal total capacity. Therefore, for the perfectly symmetric system, distributed techniques are capable of achieving the maximal information theoretic capacity. We note that in [8] this ability to achieve the global optimum using the greedy algorithm is attributed to the symmetry of the system (this analysis is based upon the Ph.D. Dissertation of Popescu [104] that shows that asymmetry causes the greedy algorithm to fail to converge for networks with more than 2 users). Therefore, despite the fact that this work proves that the optimal theoretic capacity can be achieved in a distributed way for a perfectly symmetric system, it does not tell us much about the asymmetric system except that the technique used here is not guaranteed to achieve optimal total capacity in an asymmetric system. Furthermore, perfectly symmetric systems are uncommon in practice.

6.1.3 Sample Cases

In this section, we hope to explain why higher symmetry causes a lower benefit of collaboration on average. Specifically in situations with high symmetry values, we show that selfish behavior tends to lead to near optimal results. We begin by introducing the binary channel and a few sample cases. We then explain some analytic reasoning behind why symmetry has an impact on the ability of distributed techniques to achieve near optimal performance.

Binary Channel

We use the simple binary interference channel in this section to analyze the optimality of distributed techniques. Specifically, we focus on the case in which a single user is trans-

mitting on a channel and another user senses the channel and chooses whether to transmit simultaneously on the channel. The setup of the binary interference channel can be seen in Figure 6.1.

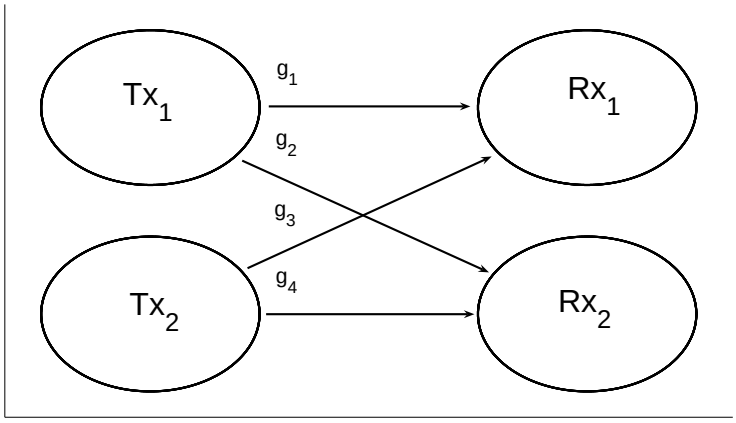

Figure 6.1: Binary Interference Channel

Because we determined previously that the situations in which a benefit to collaboration exists are situations in which interference levels tend to determine the performance, we neglect noise in these sample cases. We note that in typical situations with two users and a single channel the optimal result tends to be either that the user with the best performance is given priority to the channel or both users transmit if they can both achieve good performance simultaneously (not causing harmful interference to each other). For each sample case below, we will consider an $SIR > 3$ to be a good channel and an $SIR < 3$ to be a bad channel. Therefore, we will define optimal performance to be either both users transmit if both users can achieve good performance or the user with the best direct channel gain transmits and the other user does not if both users cannot achieve good performance simultaneously.

We note that in this model, our symmetry metric is equal to 1 when $g_1 \cdot g_3 = g_2 \cdot g_4$. For our distributed approach, we will say that, if the second user senses a good channel, the second user will transmit full power on that channel. If the second user transmitting causes the first user to have a bad channel, the first user will stop transmitting. We note that the definition of a good and bad channel, the definition of the optimal, and the distributed technique can vary but the way we have defined these values is only for demonstration purposes.

We first consider using the distributed approach in the case in which $g_1 = 1$, $g_2 = 0.1$,

$g_3 = 0.1$, and $g_4 = 1$. In this scenario, $M_{sym} = 1$. If user 1 is already present on the channel, user 2 will see a possible SIR of 10 and will also transmit. User 1 will continue to transmit because its SIR will be 10 when user 2 is transmitting. In this case, neither user interferes largely with the other user, so both users can transmit simultaneously without negatively affecting each other. In this scenario, our distributed approach achieves the optimal result. We notice that symmetry is not very important when users do not tend to interfere with each other because all users can still achieve good capacities without collaboration. This means symmetry will only be important in cases where users tend to degrade the performance of other users through interference.

Next, let us explore an asymmetric case in which $g_1 = 1$, $g_2 = 0.1$, $g_3 = 1$, and $g_4 = 0.5$. In this scenario, $M_{sym} = 0.05$. If user 1 is already present on the channel, user 2 will see a good channel and transmit also. The addition of user 2 will cause the channel of user 1 to be bad, so user 1 will stop transmitting. Similarly, if user 2 is present on the channel originally, user 1 will see a bad channel and not transmit. Our distributed approach allows user 2 to transmit but does not allow user 1 to transmit; however, the optimal performance in this scenario is for user 1 to transmit and user 2 to not transmit. The asymmetry in the system causes the distributed approach to yield a sub-optimal result.

We will now explore an interesting symmetric case where $g_1 = 1$, $g_2 = 1$, $g_3 = 0.1$, and $g_4 = 0.1$. In this scenario, $M_{sym} = 1$. If user 1 is present on the channel at first, user 2 will see a bad channel and will not transmit. However, if user 2 is present on the channel originally, user 1 will see a good channel and transmit full power. This will cause user 2 to have a bad channel and stop transmitting. We note that in both cases user 1 receives the channel and user 2 leaves. Thus the optimal result is reached because user 1 has the best direct channel gain ($g_1 = 1$ while $g_4 = 0.1$).

Symmetry Analysis

The question we hope to answer then, is why do we tend to obtain the optimal result when the channel is symmetric and what causes us to not achieve the result when the channel is asymmetric? We define a value k to be equal to both sides of our symmetry metric (i.e., $k = g_1 \cdot g_3 = g_2 \cdot g_4$). Since in the symmetric channel we know $g_2 \cdot g_4 = k$, we know that if user 2 has a good channel ($g_4 > 3 \cdot g_2$), we can expect that g_2 will be small relative to g_4. This means when user 2 has a high direct channel gain compared to other users, the interference on that channel caused by other users is expected to be small. On the contrary, if user 2 has a low direct channel gain compared to other users, the interference on that channel caused by other users is expected to be high. This means that the users with the best direct channel gains are given preference to the channel because the users with the worse direct channel gains will see higher interference power on the channel (and be more likely to not transmit).

If we extend this to the multi-channel case, a user with a high direct channel gain should be affected little by interference from other users making that user more likely to successfully transmit on a channel. On the contrary, a user with only bad channel gains will see high interference on its channels from other users, making it unlikely that this user will find a good channel. In this manner, the good channels are given to the users with the best channels.

However, this result does not extend to the asymmetric case because of cases like Sample Case 2. In this case, the user with the worse channel caused significant interference to the user with the better channel causing the user with the better channel to leave. It is in situations like this that the users with the worse channels can negatively impact the users with the better channels due to the asymmetry of the channel. Therefore in a symmetric system, users can make decisions that take into consideration their impact upon other users (despite being selfish decisions); but in asymmetric scenarios, the greedy approach makes decisions that can not consider the impact upon other users and thus are not guaranteed to be nearly optimal.

6.1.4 Linear Fit

We conclude this section on asymmetry by running simulations and plotting the benefit of collaboration against the symmetry metric and making a linear interpolation of the data points to show the impact of symmetry on the benefit of collaboration. We plot the case in which we have 4 users and 2 channels in 6.2 and the case in which we have 8 users and 2 channels in 6.3.

We notice from the individual data points that, while the linear fit shows a correlation between symmetry and the benefit of collaboration exists, our symmetry metric does not always define whether a benefit to collaboration will exist. Many cases exist where our symmetry metric would seem to suggest a low (or high) benefit of collaboration where we actually receive a high (or low) benefit of collaboration. From the linear fit, we see that the benefit of collaboration is higher as the symmetry metric becomes lower for both 4 users and 8 users. Furthermore, the benefit of collaboration exists both when we have 4 users and when we have 8 users. For the 4 user case, we expect the average benefit of collaboration to increase by about 6% for each drop of 10dB in the symmetry metric. In the 8 user case, a 10dB drop in the symmetry metric results in an expected increase of the benefit of collaboration of about 12%. We note that, while the slope of the curve in the 8 user case is higher, the gains of collaboration are also higher and the median symmetry metric is lower, so the slope being steeper does not necessary relate to a higher correlation.

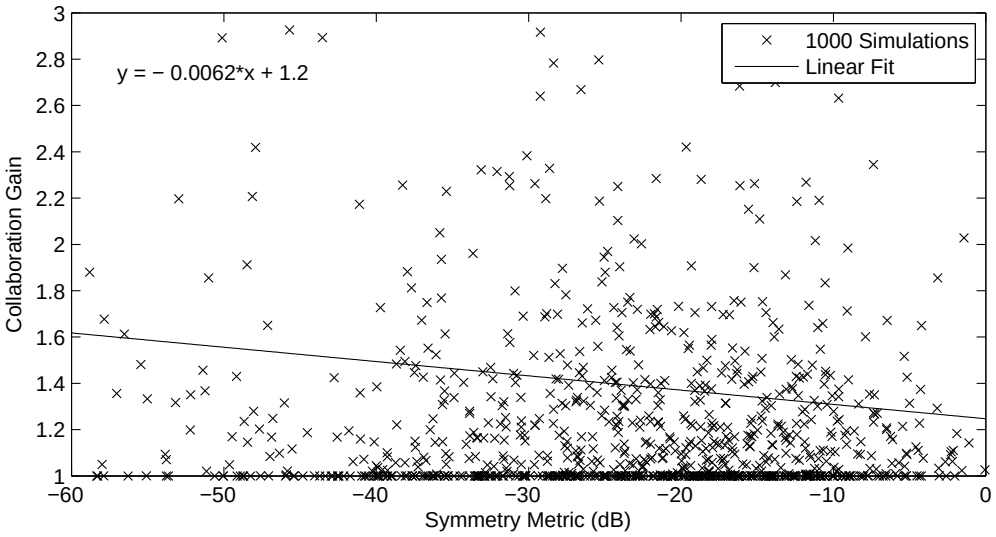

Figure 6.2: Correlation between Benefit of Collaboration and Symmetry Metric for 4 Users, 2 Channels

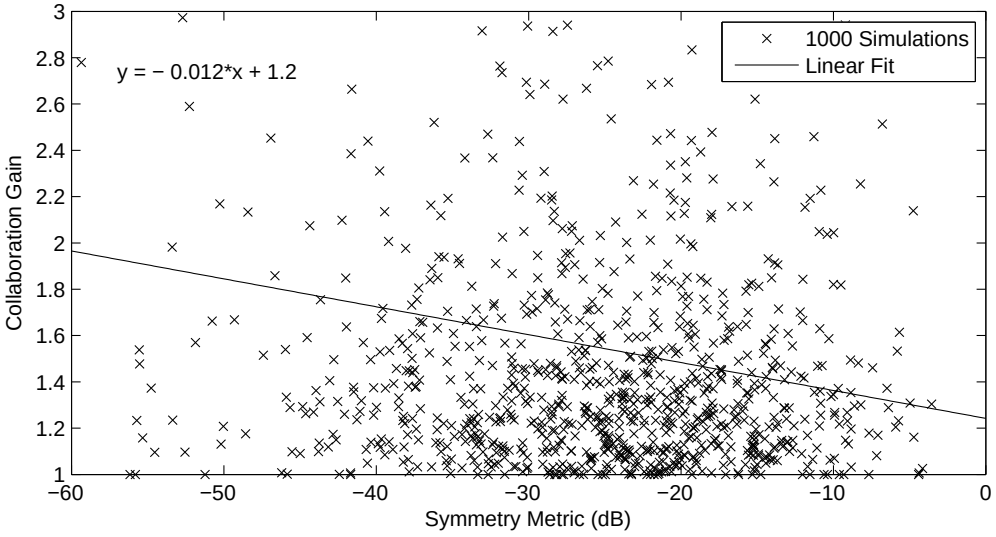

Figure 6.3: Correlation between Benefit of Collaboration and Symmetry Metric for 8 Users, 2 Channels

6.2 Dissecting the Gains of Collaboration

In this section, we dissect the collaboration gains and determine which components are responsible for the gain we find when we use collaboration. We begin by measuring the benefit of optimizing sum-rate with full power control by comparing it to optimizing sum-rate with binary power control. We then determine the value of being able to turn users off by comparing optimal sum-rate with binary power control to optimal sum-rate with no power control. In the next section, we determine the benefit of channel selection optimization by comparing the optimal sum-rate with no power control result to sensing (which inherently has no power control). We conclude this section by comparing the fairness gains for using each technique.

6.2.1 Full Power Control

In Figure 6.4, we plot the benefit of using full power control instead of binary power control for sum-rate optimization. As we can see, almost no benefit exists for full power control when optimizing sum-rate. Remember from Section 5.2 that we determined that this benefit does exist when a maximum spectral efficiency constraint is placed on the channels. This is due to the fact that, if a user achieves this maximum spectral efficiency and it increases its power further, it only increases the interference seen by other users. By using full power control, users are capable of minimizing interference power caused to other users while still achieving the maximum spectral efficiency. We also saw that a benefit exists to using full power control over binary power control when optimizing proportional fairness and when optimizing the number of admitted users given a target rate; again, this is achieved by preventing the users with the best channels from causing excessive interference to the other users.

6.2.2 Binary Power Control

In Figure 6.5, we can see the benefit of using binary power control instead of no power control when optimizing sum-rate. It is obvious from this plot that the ability of the optimization to turn users off gives a large benefit when the number of channels is high and the number of users is low (specifically when the users cannot be accommodated by the number of channels without causing excessive interference to each other). Specifically, the maximum gain in this plot exists where the number of users is highest and the number of channels is lowest, and the minimum gain exists where the number of users is lowest and the number of channels is highest. Therefore, the ability of the optimization to turn users off is a significant portion of the collaboration gain.

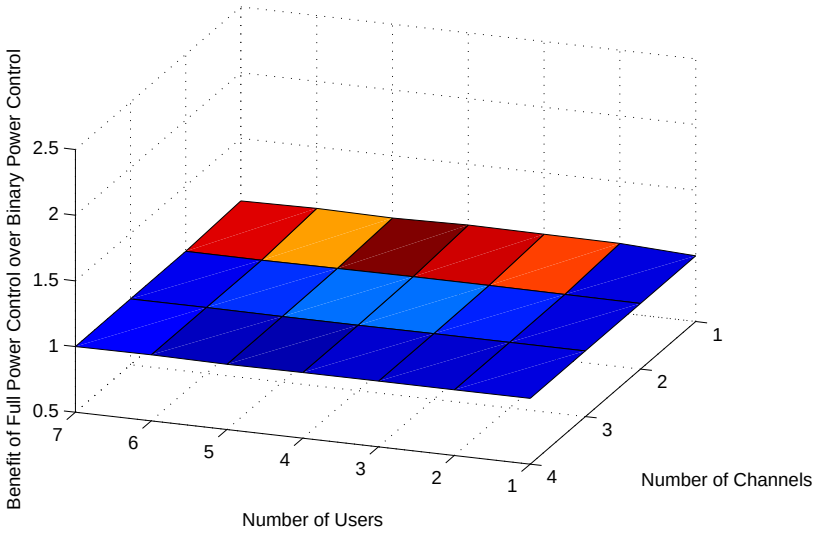

Figure 6.4: Benefit of Full Power Control vs. Binary Power Control

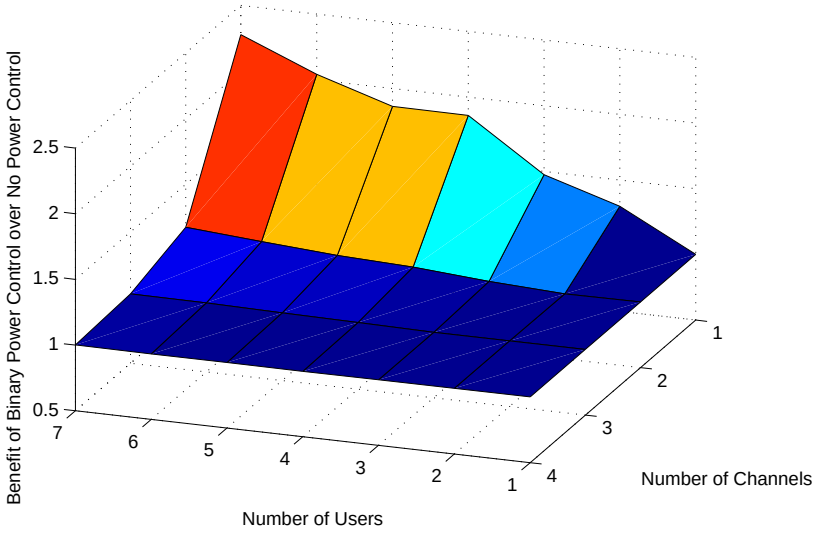

Figure 6.5: Benefit of Binary Power Control vs. No Power Control

6.2.3 Channel Selection

In Figure 6.6, we can see the benefit of using no power control when optimizing sum-rate compared to a single round of SC sensing. Since both of these situations involve no power control, we are specifically considering the impact of optimal channel selection. We notice that, when we have a single channel, no benefit exists because channel selection can play no benefit when there is only a single channel. The largest benefit of channel selection exists when two channels exists (for $N = 7$ this is around 23%), and as the number of channels increases past two, the benefit tends to slowly decrease for a fixed number of users. The decrease in channel selection gain as the number of channels increases past two is due to users spreading transmission amongst more channels causing the expected interference per channel to decrease. Since the benefit of collaboration is due to the interference avoidance, the benefit of collaboration decreases as there is less interference to avoid. As the number of users increases, the expected interference per channel increases which means that as the number of users increases, the benefit of channel selection increases. To show how multiple rounds of sensing would affect the benefit of optimal channel selection, we plot the benefit of using no power control optimizing sum-rate compared to ten rounds of SC sensing in Figure 6.7. We notice that, even with multiple rounds of sensing, the same trends occur for the benefit of optimal channel selection. Again, the maximal gain occurs with $N = 7$ and $M = 2$ yielding a benefit of around 18%.

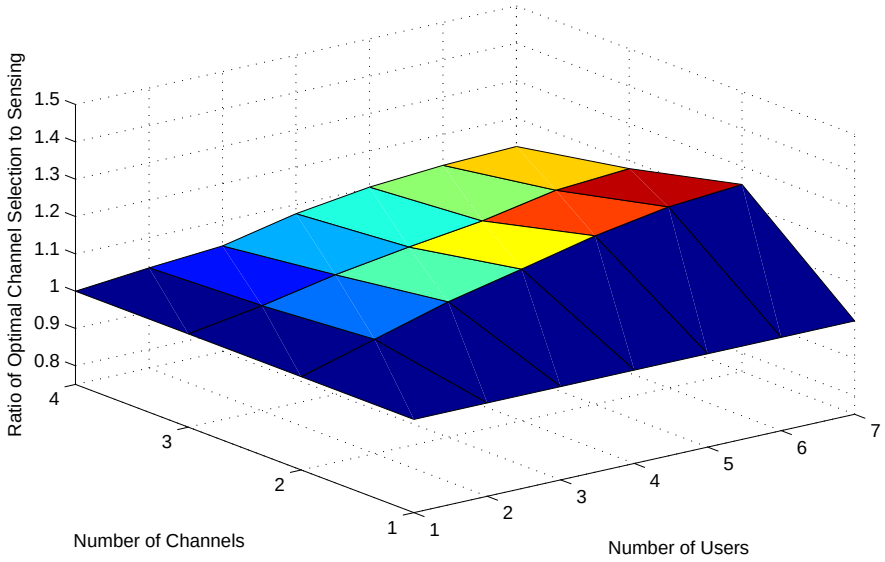

Figure 6.6: Benefit of Optimal Channel Selection over Sensing (1 Round)

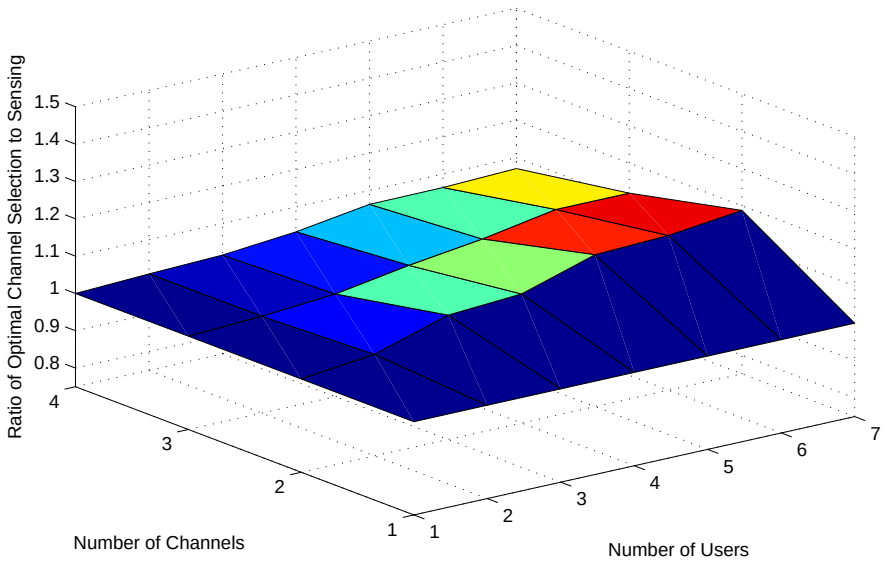

Figure 6.7: Benefit of Optimal Channel Selection over Multiple Rounds of Sensing (10 Rounds)

6.2.4 Power Control Conclusions

We have plotted the benefit of optimization with different levels of power control over sensing for a single channel in Figure 6.8. When we only have a single channel, there is no benefit to channel selection (as is expected). There is a relatively large benefit to binary power control (being able to turn users off) which increases as the number of users increases. We determined that this benefit to binary power control exists when all of the users cannot be accommodated by the available set of channels without causing excessive interference to each other. In these situations, it is beneficial to turn some users off to benefit other users. Full power control has very little benefit over binary power control in this situation as we have found is generally the case when optimizing sum-rate.

We plotted the benefit of optimization with different levels of power control over sensing for two channels in Figure 6.9. We can see that when we have two channels the benefit of channel selection (optimization with no power control over sensing) increases significantly (about 25% with 7 users). The benefit of binary power control over no power control decreases significantly because more users can be accommodated without harmful interference with the addition of the second channel. The benefit of full power control over binary power control is still minimal.

We plotted the benefit of optimization with different levels of power control over sens-

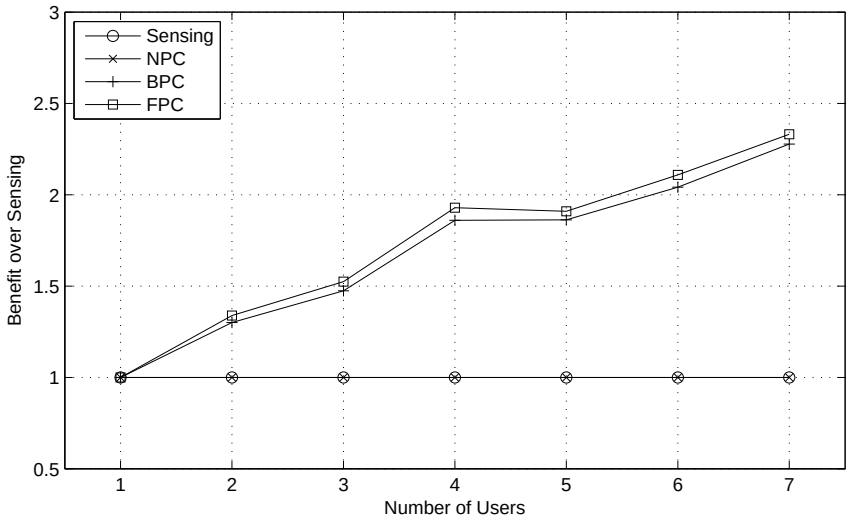

Figure 6.8: Benefit of Different Levels of Power Control over Sensing ($M = 1$, $SNR_{min} = 3dB$, and $MUI = 5$)

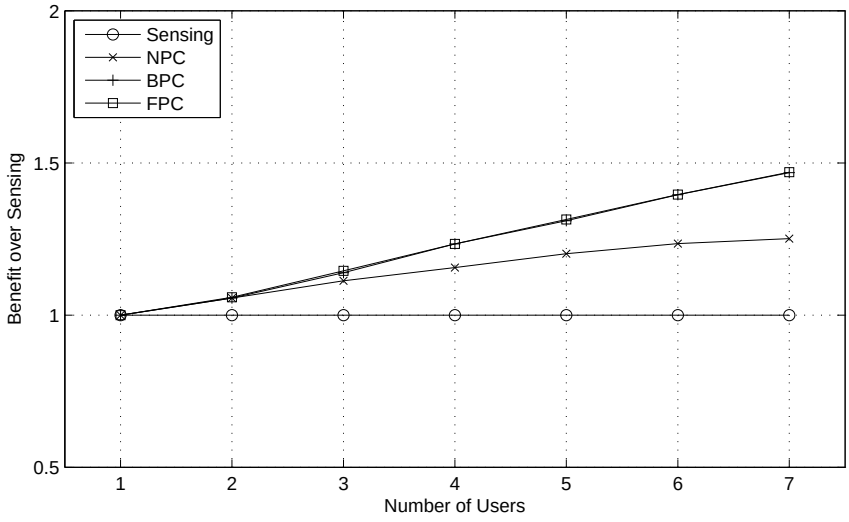

Figure 6.9: Benefit of Different Levels of Power Control over Sensing ($M = 2$, $SNR_{min} = 3dB$, and $MUI = 5$)

ing for three channels in Figure 6.10. The benefit of channel selection decreases as we add the third channel (to around 20% with 7 users). This benefit decreases because as more channels are added the interference per channel decreases, so the benefit of collaborative channel selection should be less important. The benefit of binary power control over no power control also decreases to almost zero because users tend to be better accommodated and cause less harmful interference to each other. Again, the benefit of full power control over binary power control optimization is insignificant.

We plotted the benefit of optimization with different levels of power control over sens-

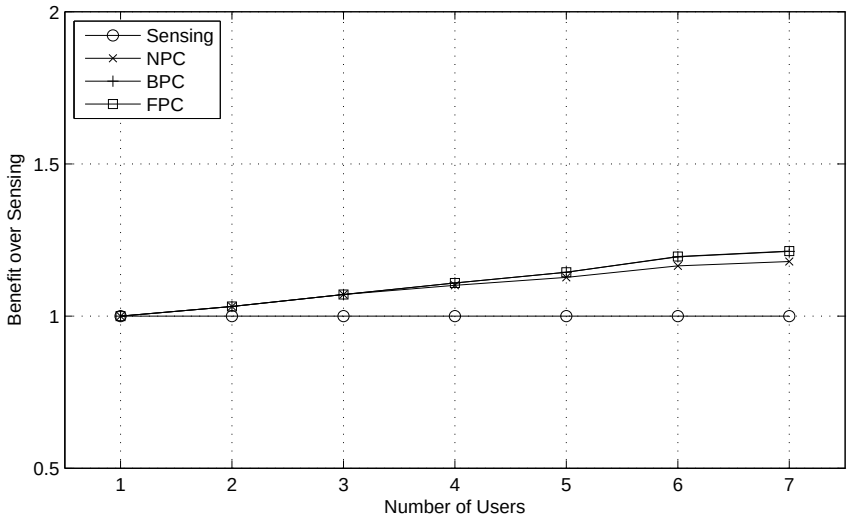

Figure 6.10: Benefit of Different Levels of Power Control over Sensing ($M = 3$, $SNR_{min} = 3dB$, and $MUI = 5$)

ing for four channels in Figure 6.11. The benefit of channel selection decreases further with the addition of the fourth channel (to around 15% with 7 users) and the benefit of binary power control over no power control decreases to nearly nothing.

6.2.5 Fairness

We now consider how different levels of power control impact the fairness. We plot the average Jain fairness of sum-rate optimization with full power control in Figure 6.12 and the average Jain fairness of sensing in Figure 6.13. When we have a single user, the Jain fairness is always equal to 1. As the number of users increases and the number of channels decreases, the fairness becomes worse. For the sum-rate optimization plot, this is due to more users being given no capacity or turned off (i.e. it is better to sacrifice some users to improve system capacity). For the sensing plot, this is due to the increasing variability among the

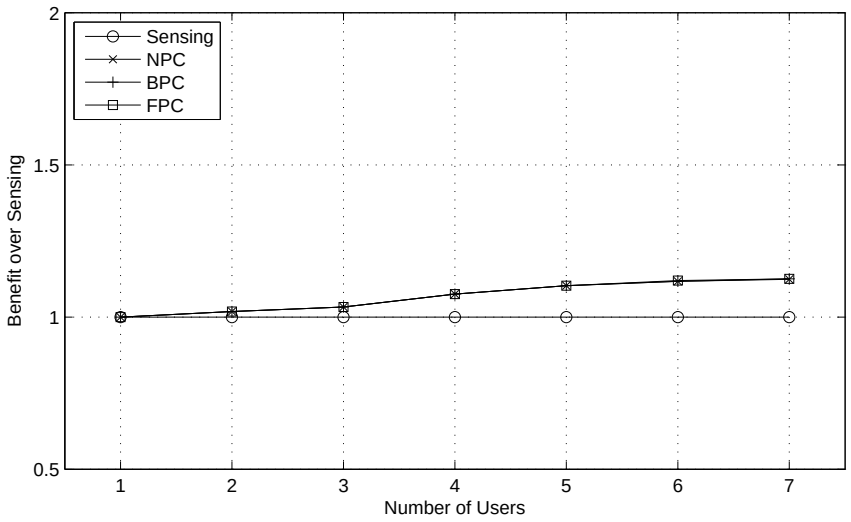

Figure 6.11: Benefit of Different Levels of Power Control over Sensing ($M = 4$, $SNR_{min} = 3dB$, and $MUI = 5$)

individual capacities of the users as the interference per channel increases (some of the best users maintain good capacities while some are degraded by the interference). We note that we can see a slight difference between the average Jain fairness values of optimization with FPC and sensing, but we will explore this difference in more detail below.

To compare the different levels of power control, we define the differential Jain fairness of technique a compared to technique b to be the expected Jain fairness of technique b subtracted from the expected Jain fairness of technique a. In Figure 6.14, we can see the differential Jain fairness of full power control sum-rate optimization compared to binary power control sum-rate optimization. We can see in the plots that the Jain fairness differential between the two techniques is nearly negligible. This is because we determined that these two techniques are nearly identical; therefore, the metrics of each technique should be nearly identical. In Figure 6.15, we can see the differential Jain fairness of binary power control sum-rate optimization compared to no power control sum-rate optimization. We can see that the fairness of binary power control is higher when the number of users is larger and the number of channels is smaller. This means that when binary power control yields a sum-rate benefit over no power control, it tends to also yield a Jain fairness benefit. In Figure 6.16, we can see the differential Jain fairness of no power control compared to sensing. We can see again that the benefit to the Jain fairness of using no power control is largely in the same region where there exists a sum-rate benefit. This benefit to the Jain fairness becomes larger as the number of users and channels increase. The main results we can see from this analysis is that, when collaboration yields gains to the sum-rate, it also tends to yield gains

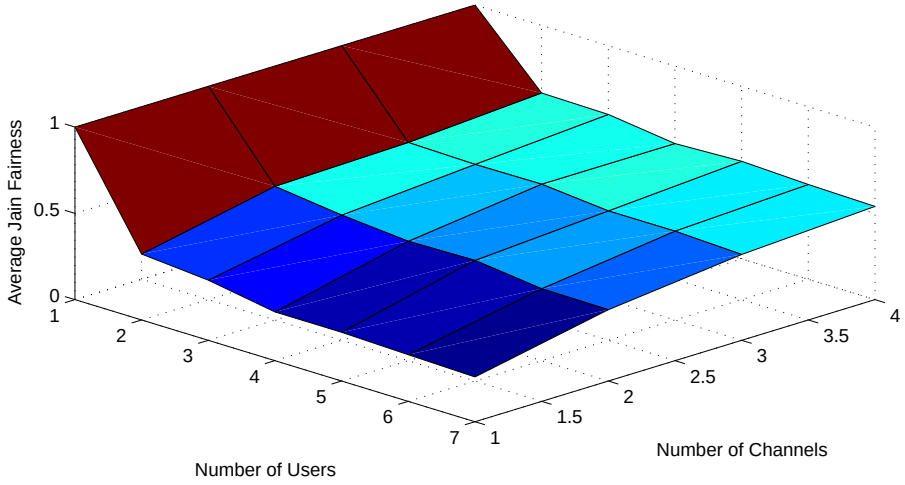

Figure 6.12: Jain Fairness of Full Power Control Sum-rate Optimization

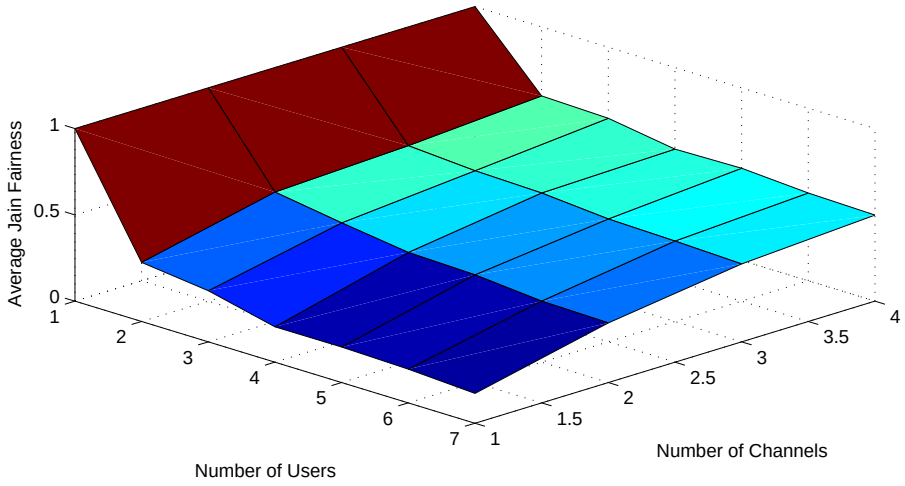

Figure 6.13: Jain Fairness of Sensing

to the fairness.

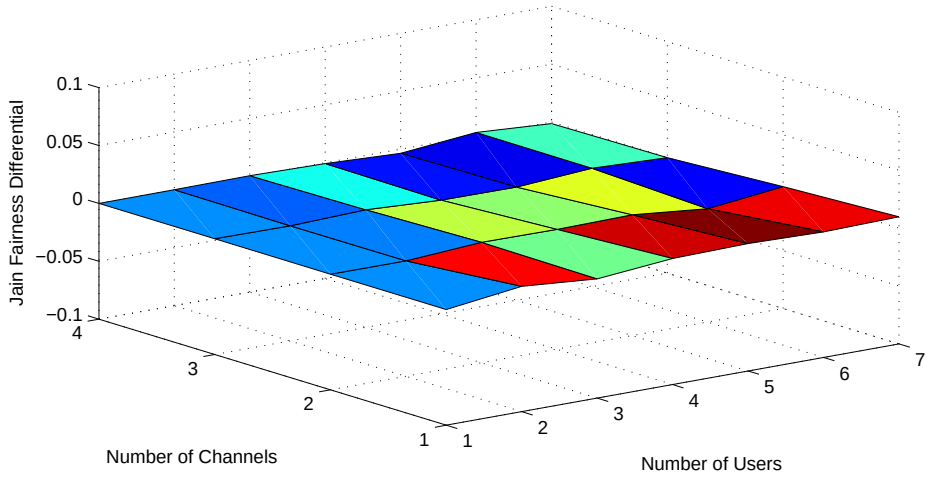

Figure 6.14: Impact on Jain Fairness of Using Full Power Control vs. Binary Power Control Sum-rate Optimization

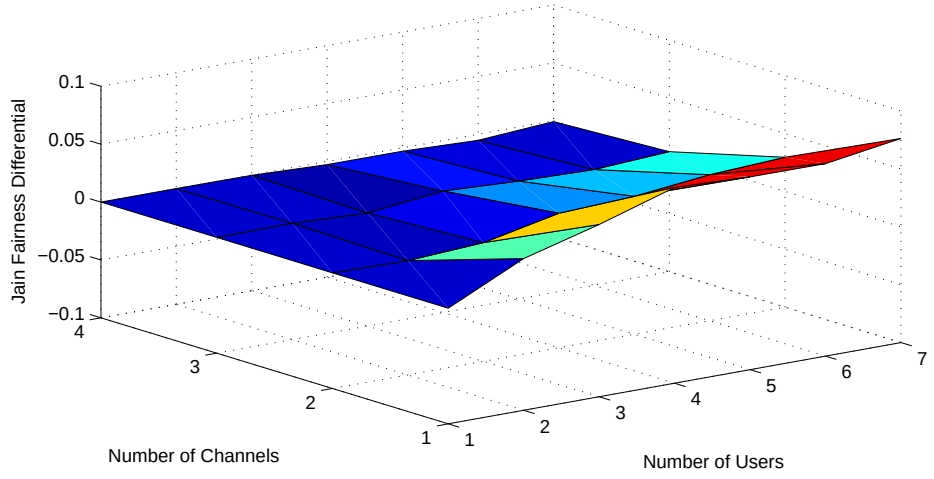

Figure 6.15: Impact on Jain Fairness of Using Binary Power Control vs. No Power Control Sum-rate Optimization

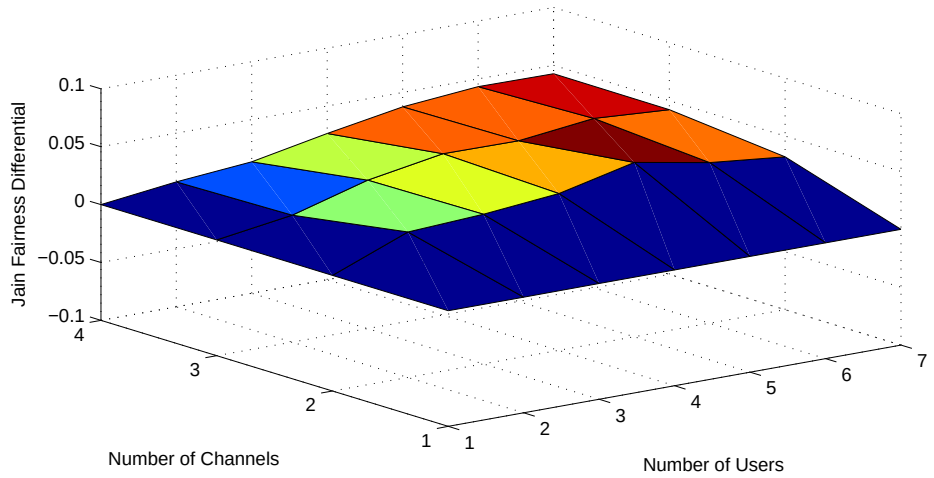

Figure 6.16: Impact on Jain Fairness of Using No Power Control Sum-rate Optimization vs. Sensing

6.3 Techniques to Achieve Fairness and Capacity Gains

We now consider how we might be able to improve the performance of distributed techniques in an attempt to achieve better performance than the sub-optimal sensing approach. Since we determined that a significant portion of the benefit of collaboration was in binary power control, we consider a case in which we use sensing, but we implement a minimum capacity that users must achieve to be allowed to transmit. Since the benefit of collaboration mainly results from interference avoidance, we also consider a time division technique where users access the channel one at a time in a contention-based (non-collaborative) approach. We conclude this section with a comparison of these techniques with several previously studied techniques.

6.3.1 Sensing with Cutoff Capacity

We recall from the previous section that a significant portion of the benefit of collaboration occurs from being able to turn users off. In this section, we develop a sensing technique that attempts to achieve some of the gains associated with binary power control by allowing users to be turned off. Specifically, we assume that if a user in unable to achieve some cutoff capacity, that the user does not transmit any power. By using this cutoff capacity, we force users that are unable to achieve at least some minimal capacity to turn off, minimizing the interference seen by the users with the best channels. In Figure 6.17, we plot the CDF of the sum-rate per user using sensing with a cutoff of 0.05 bps and 1 bps. Even with the minimal cutoff of 0.05 bps, we close the gap between sensing and SC optimization by about 25%. By using a minimal cutoff of 1 bps, we achieve about 50% of the gains that we achieve through full collaboration. We note that as we increase the cutoff, we progressively minimize interference and tend to allow the best users to achieve the maximal capacities; however, as we increase the cutoff, we also eliminate users that might be able to enhance the sum-rate of the system without causing significant interference to the best users. It is because of this second fact that we do not expect the performance of this method to fully approach the benefit that we achieve using full collaboration. In Figure 6.18, we plot the CDF of the Jain fairness for the same cases. This plot is very similar to the CDF of the sum-rate per user in that using the cutoff techniques causes the fairness to increase beyond regular sensing and become closer to the fairness of the optimization. However, we do expect that as we increase the cutoff capacity that our Jain fairness will reach a maximum and begin to decrease as more users become cutoff.

We compare the capacity benefit of cutoff techniques for two channels and varying numbers of users in Figure 6.19. In general, we can see that the benefit of a cutoff of 1 bps/Hz yields about half of the benefit of collaboration and using a cutoff of 0.05 bps/Hz yields a little less than 25% of the benefit of collaboration. We note that the loss in sum-rate of sensing with the cutoff of 1 bps/Hz when there is a single user is due to that user being turned off when it cannot achieve 1 bps/Hz which yields a slightly lower sum-rate per user on average.

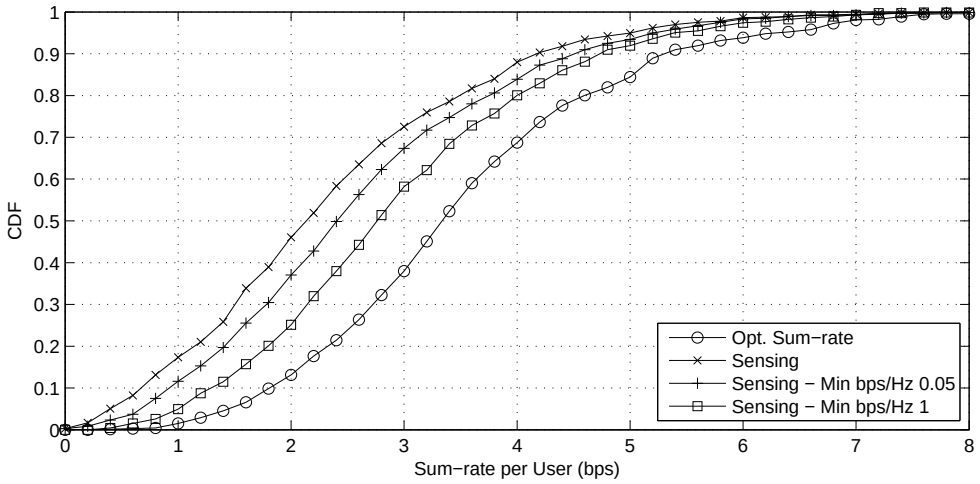

Figure 6.17: CDF of Sum-rate per User for Cutoff Methods ($N = 6$, $M = 2$, $SNR_{min} = 3dB$, and $MUI = 5$)

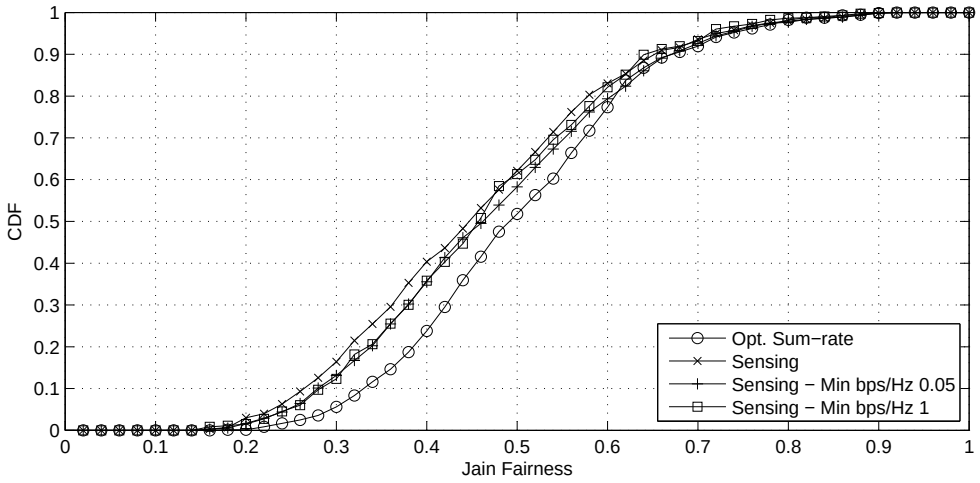

Figure 6.18: CDF of Jain Fairness for Cutoff Methods ($N = 6$, $M = 2$, $SNR_{min} = 3dB$, and $MUI = 5$)

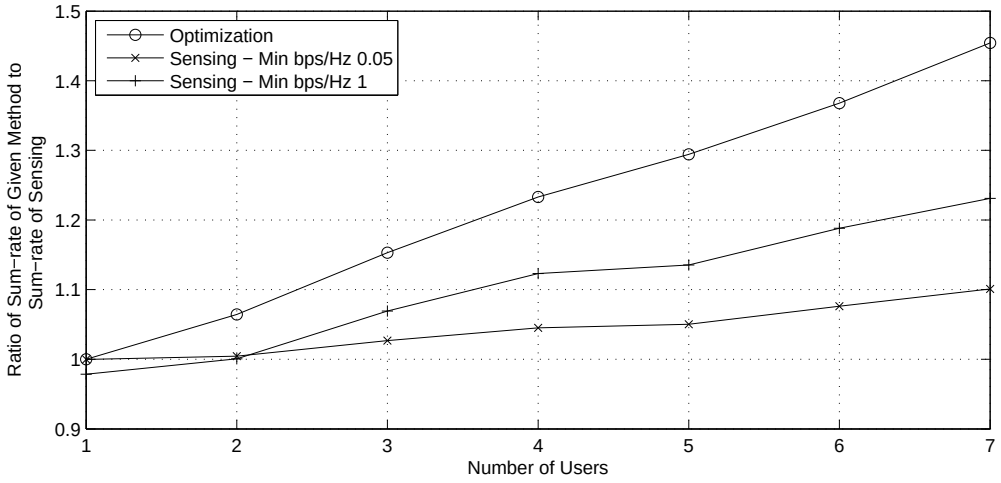

Figure 6.19: Benefit of using Cutoff Methods for Sum-rate per User ($M = 2$, $SNR_{min} = 3dB$, and $MUI = 5$)

6.3.2 Time Division Techniques

The second distributed technique that we could use to achieve the collaboration gain through interference avoidance is a time division technique. Using time division allows users to transmit without multi-user interference; however, time division techniques are sub-optimal because they do not consider that two users may be able to simultaneously transmit on the same channel without causing harmful interference to each other. Because the purpose of introducing these techniques is to avoid collaboration, it is necessary to use contention based techniques for random channel access; however, for the simulations that we perform in this section, we will assume that the contention based techniques require no overhead.

We consider the impact of giving preference (a.k.a., more time) to the worst or best users by introducing a weighting factor, κ. Determination on which users have the best capacities is based upon the SNR capacity for each user i, $C_{SNR,i}$ (this is the capacity a user can achieve assuming no interference). Based upon these achievable capacities, if $\kappa > 0$ preference to the channel is given to those users with the best achievable capacities, if $\kappa < 0$ preference to the channel is given to those users with the worst achievable capacities, and if $\kappa = 0$ all users are given the same preference (or equal time) on the channel. We introduce two techniques to implement the weighting factor, but the main goal of the weighting factor is to cause the fraction of time that user i transmits on the channel to be approximately,

$$\frac{(C_{SNR,i})^{\kappa}}{\sum_{i=1}^{N}(C_{SNR,1})^{\kappa}} \tag{6.2}$$

The first technique we introduce to weight the users access times is to use a variable packet length (length is in units of time). In this technique, all users are equally likely to gain access to the channel when the channel becomes inactive; however, the packet length (in time) a user sends is proportional to $(C_{SNR,i})^{\kappa}$. For example, if we use $\kappa = 1$ the users with better capacities will be allowed to transmit for longer packet durations. We note that this technique requires that all users agree to a correlation between packet length and achievable capacities ahead of time, but this could be based on regulation and thus does not require collaboration (i.e., we could set $\kappa = 0.5$ and set a packet length for $C_{SNR,i} = 1bps/Hz$ to be 20 ms; from these values packet lengths can be derived for all other values of $C_{SNR,i}$).

The other alternative to weight the users access times involves using variable window sizes. Unlike the previous technique, this technique gives preference by changing the probability a user receives access to the channel but keeps the packet length constant. In most random access techniques, windows are used to avoid collisions. In these techniques, when a channel becomes idle a user desiring to transmit on the channel takes its window size (a fixed length of time) and randomly selects a point of time within the window. The user waits until that random length of time has passed, and if the channel is still idle after waiting, the user begins to transmit. Windowing in random access is used to avoid the situation in which multiple users see the channel become available and transmit simultaneously causing a collision. By using windowing techniques (and causing users to select a random waiting time), it is likely that one of these two users will begin to transmit before the other user avoiding the collision. By making the window length of user 1 smaller than the window length of user 2, we can make it more likely that user 1 transmits first (and receives access to the channel. Therefore, we weight the window length proportional to $(C_{SNR,i})^{-\kappa}$. Again this requires all users to have an agreed upon value for κ and an agreed upon relationship between $C_{SNR,i}$ and window length, but this does not require collaboration.

In Figure 6.20, we can see the CDF of sum-rate per user using the time division techniques as compared to sensing and optimal sum-rate. We can see that giving all users equal time on the channel ($\kappa = 0$) yields the worst performance. Using a proportional packet size ($\kappa = 1$) or a proportional window size ($\kappa = 1$) yields nearly equivalent performance. The average sum-rate per user of both of these are slightly less than that of sensing, but the time division techniques more consistently yield performance around the average sum-rate per user (the variance of sum-rate per user for time division techniques is smaller). From this, we can see that using time-division techniques, even those that favor the best users, do not tend to achieve near optimal sum-rate. However, we notice that, unlike sensing, we can achieve a balance between sum-rate and fairness by varying our weighting factor κ. For example, by using $\kappa = -1$ (a packet size that is inversely proportional to the achievable rate), we achieve nearly perfect fairness as this causes each packet to have the same amount of data (it is only nearly perfect because power constraints limit how much of the time a user may be on the channel). For example, when we use multiple channels (M), each user can only receive a

time proportional to one over the number of channels ($\frac{1}{M}$) on each channel due to power constraints. We plot the CDF of the Jain fairness for the same cases in Figure 6.21. We can see that the fairness of the equal time approach far exceeds the fairness of all the other approaches. Prioritizing the best users with $\kappa = 1$ causes the fairness of the time division approaches to be slightly worse than the fairness of sensing. We can see from these plots that the value of κ allows us to achieve a balance between sum-rate and fairness. From the equal time ($\kappa = 0$) we get a significantly lower sum-rate but a significantly higher fairness; but from prioritizing the best users ($\kappa = 1$) we get a higher sum-rate but a lower fairness. We expect with $\kappa = -1$ that we can increase our fairness even further, but this will come at the cost of decreasing sum-rate.

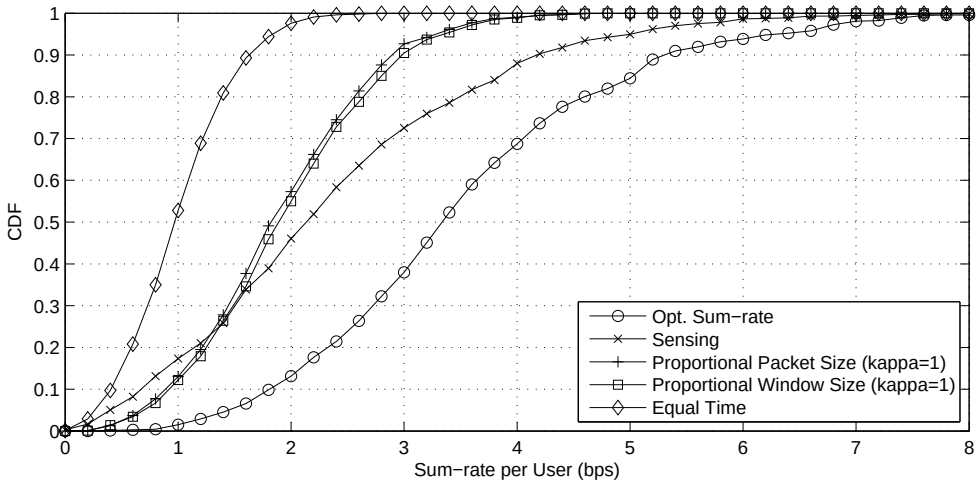

Figure 6.20: CDF of Sum-rate per User for Time Division Methods ($N = 6$, $M = 2$, $SNR_{min} = 3dB$, and $MUI = 5$)

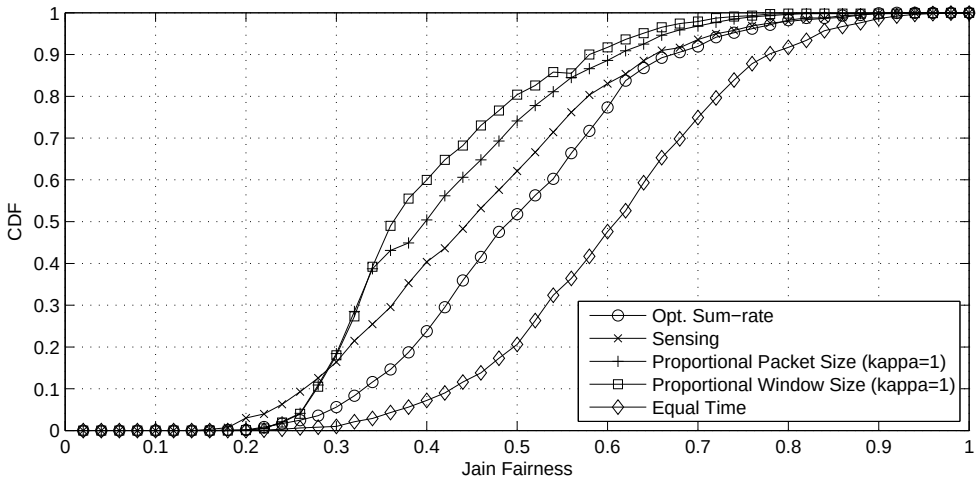

Figure 6.21: CDF of Jain Fairness for Time Division Methods ($N = 6$, $M = 2$, $SNR_{min} = 3dB$, and $MUI = 5$)

6.3.3 Comparison of Techniques

In this section, we compare the techniques that we have introduced in this section with some of the optimal results we considered throughout this report. In Figure 6.22, we compare the optimal collaborative results with sensing with a cutoff capacity and with time division with variable packet length and variable contention window sizes for $N = 6$ and $M = 2$. We compare the sum-rate per user and Jain fairness of each of these techniques simultaneously. We notice immediately the balance between fairness and capacity that we can achieve through using the time division approaches by using different proportion values (κ. However, the sum-rate that these time division approaches can achieve is not higher than the sum-rate achievable by sensing. We note that the time division technique does not achieve perfect fairness when the packet size is inversely proportional to the achievable capacity as would be expected. This is due to the limit that each user can only occupy each of the M channels for a time equal to $\frac{1}{M}$ due to power limitations. This limit causes the worst users to sometimes be unable to achieve equal capacity with all other users, creating the inability to achieve perfect fairness. We notice that the sensing with the cutoff of 0.05 causes the capacity and fairness to increase toward the optimal sum-rate with full collaboration. However, as the cutoff capacity increases, the fairness starts to decrease. Furthermore, when the cutoff capacity reaches a value of about 1, the sum-rate starts to decrease too. Both of these effects are due to the fact that as the cutoff capacity increases, more users are cutoff.

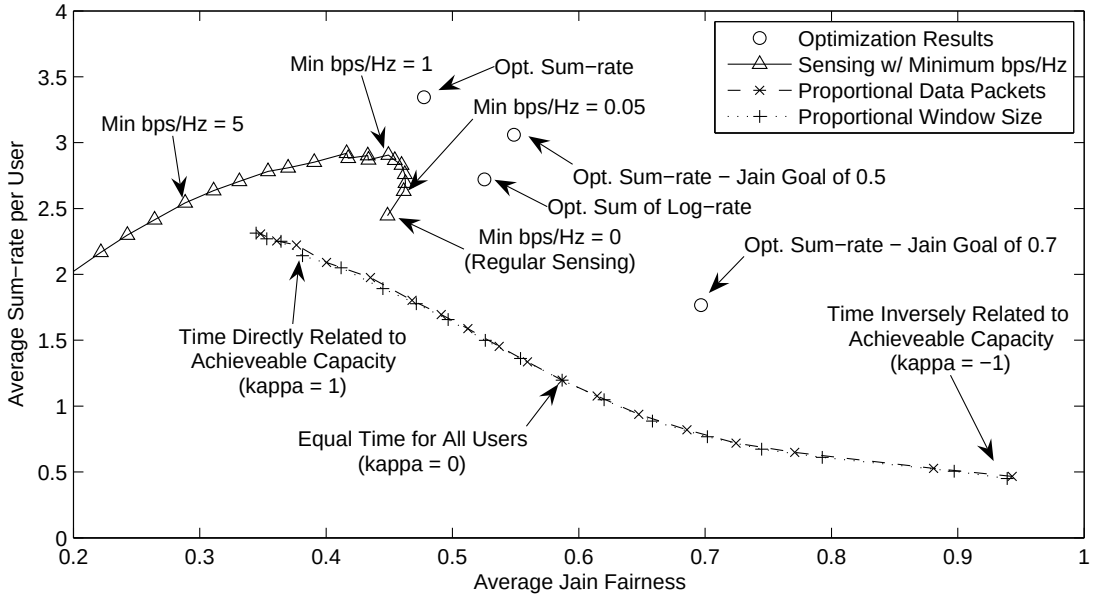

Figure 6.22: Comparison of Distributed Techniques ($N = 6$, $M = 2$, $MUI = 5$, $SNR_{min} = 3dB$)

6.4 How much Sharing / Collaboration is Needed

We now consider for each of the approaches we introduced how much sharing and collaboration is necessary. We note that we will assume for the sensing approaches that any cutoff capacity will be a value agreed upon by cognitive users ahead of time (as a regulation instead of a value decided upon). This means that the sensing approaches require no collaboration since they simply involve sensing of the channel and making decisions based upon the results of the sensing. For time division techniques, some negotiation is necessary to be able to share the channels efficiently and determine who has permission to access which channel. This negotiation overhead will cause a capacity loss as compared to the optimal results we showed in the previous section. To determine what we expect this loss to be, we spend the remainder of this section reviewing time division techniques used to share a channel efficiently.

In the class of time division techniques that assume a maximum of a single user active at any given time, there are a few major divisions. The most important division is between collaboration based and contention based. With collaboration based techniques, users negotiate the use of channels with the negotiation process occurring in another frequency band or a timeslot. These collaboration techniques have the overhead of the collaboration stage. The contention based techniques instead assume a random access technique to access the channel and allows the possibility that collisions will happen between neighboring nodes if they attempt to access the channel simultaneously. The second major division between time division techniques is whether time is slotted or not. Most collaboration based techniques assume slotted time (as it is necessary for the time slot collaboration frame), while contention based techniques tend to allow both slotted and non-slotted techniques.

The fact that the time-slotted approach requires that all users be synchronous raises complexity issues. It has been suggested in [105] that the time-slotting approach in [106] may be practical in a distributed network; however, this still requires users to collaborate. In [40], it is suggested that synchronization could be achieved through GPS; however, this is also not a good option because the power requirements to receive GPS signals is high. Since time-slotting in distributed systems seems implausible for our system model, we assume a non-synchronized non-slotted approach to channel access.

Because channel sensing is necessary in cognitive radio, we assume that for all techniques channel estimation is necessary (and therefore this overhead is the same for all collaborative and non-collaborative techniques). Contention based techniques without slotted time have an overhead that includes losses due to the improper channel sensing, the time of the contention frame, and collisions. Contention based techniques are required to adequately determine whether other nodes are currently transmitting data on the channel to prevent

from causing interference to other nodes. This channel sensing is made more difficult when the cases of the hidden and exposed nodes are considered. During the contention frame, the channel becomes idle creating a missed throughput opportunity. When many users are present, collisions may become common if the access scheme is not designed properly, and this can lead to large loss in achievable throughput on the channel.

Due to the losses associated with the time-division techniques, the capacity achieved by our optimal scheduling in section 6.4 is an optimistic estimate of the real achievable capacity of the channel. Because cognitive radios are required to be able to sense the channel to high resolution to avoid interfering with primary users, the channel sensing is not a significant factor in the consideration of the losses due to using time-division techniques. To address the hidden and exposed node problem, we introduce a technique based upon the busy burst approach in the introduction (in [40]). This technique was able to avoid the issues involved with the hidden and exposed nodes, but had an overhead of about 10%. While this paper is based upon a time slotted approach, the technique seems plausible for a non-slotted approach.

Because we will not have empty slots and we assume the overhead due to avoiding the hidden and exposed node problem is 10%, the main losses we should obtain are due to collisions and contention frames. Because the contention frame size and probability of collision are reliant upon exact system parameters, the loss due to using a collision mitigation technique may not be easily calculable. Further, using variable sized contention windows may yield different collision probabilities than if all users were to have the same window size, so analyses based upon the assumption that all users window sizes are the same are hard to translate to our system. To create an estimate of the losses due to collision, we will explore many of the common channel access techniques and determine the loss due to collision for each of them. By comparing each of these techniques, we should be able to make a reasonable estimate as to the loss due to contention for our system.

In general, two different types of communications are transmitted over wireless networks, delay tolerant and delay intolerant. In delay tolerant applications, maximizing sum-rate tends to be common practice; but in delay intolerant applications, either the number of admitted users or the minimum $SINR$ tends to be the optimization function. In this work, we focus on maximizing the sum-rate of a network which is equivalent with saying that we are focusing mainly on delay tolerant applications.

Two different types of plots are generally used to show the performance of a medium access control technique, namely, throughput-delay and throughput vs. offered throughput plots. Throughput-delay plots tend to determine what delay occurs between when a user wishes to transmit a packet and when that user actually transmits its packet given a specific loading of the channel (loading is generally given as a fraction of rate offered divided by maximum rate

of the channel). Since we focus mainly on maximizing throughput, this plot is minimally significant to us since we do not want an infinite delay, but we do not mind if the data has some delay. On the other hand, the throughput vs. offered throughput plot is largely significant to us because this tells us what percent efficiency can be reached in our system.

We note that we use estimates for our work because exact efficiencies are combinations of many variables. For example, as the packet size increases, the number of contention frames per second decreases causing the expected number of contention frames and collisions per second to decrease. Also, as the number of users increases, the expected number of collisions also increases. The number of collisions is also reliant on the expected load of the channel. For example, if a user only wishes to transmit for a fraction of the time collisions will be a lot less likely than if this user had an infinite queue and wanted to transmit continuously. It is shown in [107] that the propagation delay can have a significant impact on the efficiency of a system (efficiency ranges from about 20-100% as the propagation delay ranges from the packet transmission time to zero). Furthermore, the variable initial window size or variable packet size is likely to change the expected number of contention frames and collisions compared to a system with a fixed window size and fixed packet size. For all these reasons, it is impractical for us to derive an exact metric for the percentage of inefficiency of this system.

Compared with other research in this area that assumes all users can achieve an equal rate on the channel, we will instead assume users can achieve whatever rate the SNR on that channel allows them to obtain. Therefore, to measure the efficiency of the channel, we will measure the fraction of time the channel is used for data divided by the total time the channel could have been used. For example, if 2 seconds are devoted to avoiding collisions and 8 seconds are used to transmit data (in a 10 second time period), our channel efficiency fraction would be 0.8. We compare the channel efficiency fraction of other schemes to estimate the efficiency that could possible be obtained.

The first major MAC scheme introduced around 1970 in [108] was ALOHA. In this work, it is shown that the maximum efficiency that ALOHA can achieve is about 0.186. Non-persistent and persistent CSMA were studied in [107]. It was found that the throughput efficiency depended highly upon the ratio of propagation time to packet transmission time. For example, for non-persistent CSMA at a ratio of 1, efficiency was around 15%; at a ratio of 0.1 efficiency was around 55%; at a ratio of 0.01, efficiency was around 80%; and at a ratio of 0.001, the efficiency was around 90%. This work assumes a propagation time to packet transmission time ratio of 0.01 when this ratio is a constant; therefore, we use 0.01 to be a realistic estimate of this fraction. This means that the maximum efficiency possible is around 80% and this is only achieved when the loading rate is about 10 times the channel capacity. It is also shown that for $n-$persistent CSMA the maximum throughput occurs at different values of the offered load for different values of n. By using persistent CSMA, a decrease in the necessary loading rate to around 1 decreases the efficiency to around 60%.

Two further popular techniques that introduce a request-to-send and clear-to-send sequence before transmitting are the MACA and MACAW techniques, however these techniques do not perform carrier sense [109]. DFWMAC, the MAC used in the IEEE 802.11 standard, extends these protocols by adding the carrier sensing. Performance of DFWMAC is explored in [97]. This study emphasizes the reliance on the number to the number of bytes per packet (bpp) (64 bpp yields 33%; 256 bpp yields 66%; 1024 bpp yields 88%). Unlike the CSMA techniques mentioned above, these efficiencies to not require heavy loads in the network and do not perform poorly under heavy loads.

From this study of previous work, we can see that the DFWMAC is the most sophisticated and tends to yield the best performance. Since as we mentioned previously we are working mainly with a delay tolerant system, we can likely increase the packet size to yield good performance. We will thus assume that we can achieve the equivalent of 1024 bytes per packet (remember that we share the channel by time and thus do not have a fixed bytes per packet) and achieve an 88% efficiency of channel usage. Since it is mentioned in [97] that the DFWMAC does not completely correct the hidden node problem, we will assume the additional overhead of the busy burst technique mentioned previously in [40] which is estimated to have an overhead of around 10%. Therefore, we will assume that we can achieve around 80% efficiency using a random access scheme that avoids the hidden and exposed node issues. While this does not consider improper channel sensing, as we mentioned previously, cognitive radios are required to be highly efficient at channel sensing; therefore, we will neglect this loss.

6.5 Impact of Time on this Work

We now wish to consider the impact of burstiness on each of the techniques we study. Specifically, if users are only active for a fraction of the time, what impact will this have on each technique? We consider this for two different cases. First, we consider what impact the burstiness has on time division techniques (we do this through literature review as contention based techniques are well studied). Second, we consider what impact this has on the techniques that treat interference as noise. The specific case we consider for this second scenario is what loss we will have if we assume that all users are active but only a fraction of the users are actually active. When not all users are active, it is possible that re-optimization would be able to achieve better results for the users that are still active. We study this second technique both for sensing and optimizing sum-rate.

6.5.1 Burstiness in Time Division Techniques

In this section, we study the impact of burstiness on the time division techniques we introduced in Section 6.3. By burstiness, we mean that users do not have a queue of data waiting to be transferred that has an infinite length. In time division techniques, the consideration of burstiness is often explored with presentation of a throughput vs. offered throughput. For our analysis, this plot translates to an amount of time used for data transmission on the channel plotted against the amount of time necessary to accommodate all the data that users desire to transmit (we translate this because typically using this plot assumes that all users achieve the same rate which we do not assume). In Figure 16 in [97], we can see the throughput vs. offered throughput plot for DFWMAC. This plot shows us that when the offered load is over 100% the throughput is relatively constant. It also shows that the throughput vs. offered throughput tends to follow a linear ramp until it reaches close to the throughput percentage upper limit at which point it becomes flat. This means that users should be able to achieve good performance with anywhere from light loads (small expected fraction of active users) to heavy loads (high expected fraction of active users). Therefore, time division techniques are well suited for bursty users since it allows light loads with good efficiency and heavy loads do not detract from performance.

6.5.2 Burstiness in Techniques that Treat Interference as Noise

We now consider the impact of burstiness on the techniques that treat interference as noise and allow multiple users to access a channel simultaneously. For the purpose of simulating bursty transmission, we assume that time is slotted and that each user determines independently from slot to slot whether to transmit. The expected fraction of the timeslots that users transmit in is a value that we call the fraction of active users.

We compare the case in which a single optimization is performed across all users and when an optimization is performed during each time slot for only the active users. For the single optimization case, we assume that the optimal strategy is followed by all users regardless of whether all users are accessing the channel. For example, for the single optimization case if users 1 and 2 can each achieve 1 bps on channel 1 when transmitting simultaneously and user 2 is not active for a time slot, user 1 will still assume the presence of user 2 and only achieve 1 bps although user 1 may be able to achieve better capacity due to the reduced interference from user 2 being inactive. For the optimization being performed each time slot, we determine what users are active for each time slot and optimize over that set of users. Since this technique achieves the optimal results, we expect this technique to always outperform the single optimization.

In Figure 6.23, we plot the sum-rate per user for the active users for sensing with $N = 6$ and

$M = 2$. By saying the sum-rate per user for the active users, we mean the sum-rate achieved divided by the number of users active on the channel during a given timeslot. By defining sum-rate per user this way, we keep the expected sum-rate per user of the single optimization constant which allows us to see the benefit of multiple optimizations more clearly. We can see that the sum-rate per user of the active users decreases almost linearly as the fraction of active users increases. We note that the value of the sum-rate per user as the fraction of active users approaches zero is equal to the average sum-rate per user of one user on 2 channels. This is because as the fraction approaches zero, the large majority of the time that any user is active, only a single user is active. As the fraction of active users approaches one, the sum-rate per user approaches the average sum-rate per user for sensing with no burstiness (since burstiness becomes less relevant as users are inactive less often).

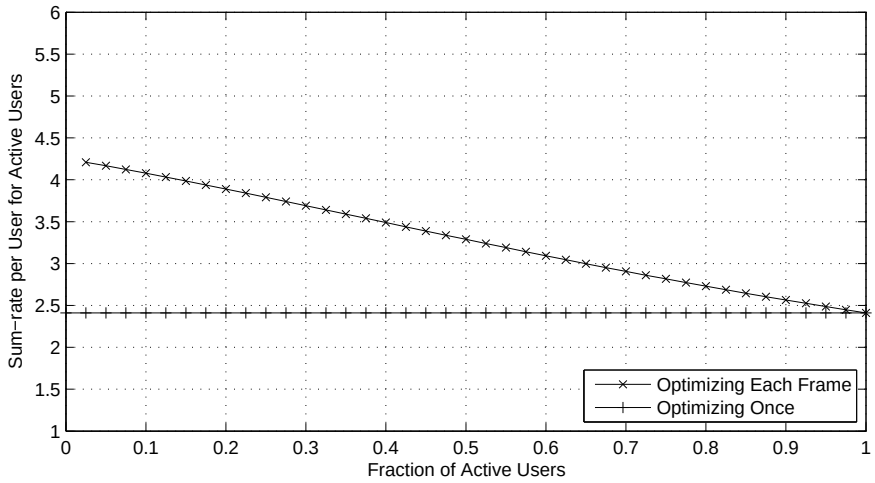

Figure 6.23: Impact of Burstiness on Performance of Sensing

In Figure 6.24, we plot the sum-rate per user for the active users for optimal sum-rate with $N = 6$ and $M = 2$. We see the same result as in the sensing plot where a linear decrease exists as the fraction of active users increases. Similar to the previous plot, as the fraction of active users approaches zero, we approach the performance of a single user with 2 channels.

From these simulations, we can see the the gain of using multiple optimizations over a single optimization has a linear relation to the fraction of active users. The maximal gain occurs as the fraction of active users approaches 0 and this performance is nearly equivalent to the performance as if a single user existed with the same number of channels. Because this performance is nearly equivalent to the performance with a single user, the optimal result with sensing and optimal sum-rate both approach the same value (optimal performance

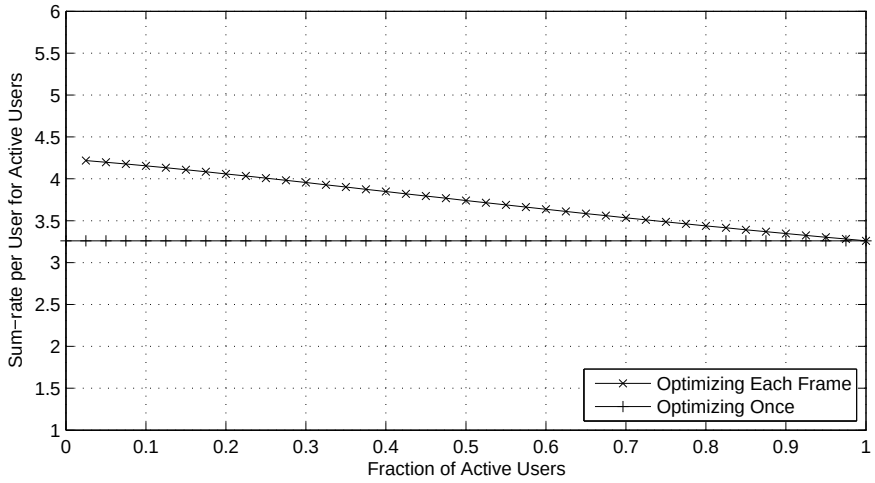

Figure 6.24: Impact of Burstiness on Performance of SC Optimal Sum-rate

for all users in a single user system is equivalent with optimal performance for each user). Considering that reoptimization in each time frame is simpler for sensing, we notice that when the fraction of bursty users is below 0.5 that the performance of sensing with reoptimization in each time slot outperforms the single optimization where we optimize sum-rate. If we assume that reoptimization in each time slot is plausible for both approaches, then we notice that the benefit of collaboration decreases as the fraction of active users decreases (this is due to fewer users being present on average in each time slot).

6.6 Conclusions

The conclusions we have made in this chapter include:

- Perfect symmetry of channel gains allows the maximal information theoretic capacity to be achieved by a distributed approach. Specifically, perfect symmetry allows the selfish motives of individual users to consider the impact on other users. However, perfectly symmetric channel gains are not common in distributed systems.
- A correlation exists between the benefit of collaboration and the value of the *symmetry metric* that we introduced, but this correlation is weak as significant outliers exist.
- Very little benefit exists to optimizing with full power control vs. binary power control.
- Optimizing with binary power control has a significant benefit over optimizing with no power control when users tend to cause harmful interference to each other. Therefore,

as the number of users increases and the number of channels decreases, the benefit to being able to turn users off increases.

- The benefit of optimal channel selection over sensing increases as the number of users increases. With a single channel, there is no benefit (since no channel selection is possible). As the number of channels increases past two, the benefit of optimal channel selection over sensing decreases due to the ability of sensing to be more likely to achieve the optimal channel selection.

- We can improve the performance of sensing by turning off users that cannot achieve at least some minimal capacity. By using a minimum bps/Hz of 0.05, we were able to achieve a little less than 25% of the sum-rate per user gain achieved by full collaboration. By using a minimum bps/Hz of 1, we were able to achieve about half of the sum-rate per user gain achieved by full collaboration.

- By using time division techniques, we can create rules that allow us to balance capacity and fairness. By highly prioritizing the best users, time division techniques achieve slightly less than the capacity achievable by sensing.

- We estimated the overhead involved with random-access techniques to be around 20%.

- When the expected fraction of active users is nearly equal to 0, the average sum-rate per user is equivalent to the average sum-rate per user for a single user (on the same number of channels). When the expected fraction of active users is equal to 1, the average sum-rate per user is equal to the sum-rate per user of all users on the same number of channels. The average sum-rate per user is linearly correlated with the fraction of active users for all values between 0 and 1. As the fraction of active users decreases, the collaboration gain decreases because the expected number of active users decreases and thus the expected interference decreases.

Chapter 7

Conclusion

Our main goal in this work was to characterize the benefit of using collaborative techniques over the non-collaborative technique which we called sensing (a technique in which each user selects the channel that gives it best performance and transmits full power on that channel to greedily maximize performance). We determined that the benefit of collaboration resulted from interference mitigation. This means that in scenarios where interference power limits capacity the benefit of collaboration is higher than when the noise power is limiting capacity. As we increase the number of users in the system, the interference power increases which causes a larger benefit to collaboration; as we increase the number of channels available in the system, the interference tends to be spread amongst more channels, causing the benefit of collaboration to decrease. We found that adding even a small amount of interference from users outside the collaboration degrades performance quickly. This has a similar effect to increasing the expected noise power since this interference power cannot be controlled, and as the expected noise power increases the benefit of collaboration decreases. Partial collaboration tends to have performance between that of full collaboration and no collaboration. The benefit of partial collaboration over no collaboration tends to be linearly related to the number of groups that users are split up into. We found that optimization can provide better Jain fairness than sensing, but with optimization the worst users tend to receive no capacity while with sensing all users receive some capacity (although that capacity may be very small). We found that, when using a system with a target rate, we still have a benefit to collaboration.

For sensing, we found that by using multi-channel sensing we can improve the performance over single channel sensing; however, this increase in performance tends to come at a cost of a decrease in fairness. By using multiple rounds of sensing, we are able to achieve slightly better capacities and slightly better fairness by increasing the capacities of the worst users, but these increases are minimal. We explored the achievable capacity and fairness when using sensing in a target-based system. We found that by increasing the target rate the capacity per user tends to increase to a maximum at which point the capacity per user de-

creases as the target rate increases. It increases capacity because it allows the best users to achieve higher capacities even though it reduces the number of users which can achieve the target. Increasing the target rate decreases the fairness because fewer users are expected to achieve the target rate as the target rate increases.

When optimizing sum-rate, very little benefit exists to using full power control as compared to binary power control due to the fact that most users either transmit full power or no power. We proved that for the 2-user case, the optimal sum-rate when using full power control is equivalent to the optimal sum-rate when using binary power control. When we introduce a maximum spectral efficiency, a benefit to full power control exists over binary power control because users will transmit the minimum power that allows them to achieve the maximum spectral efficiency to minimize interference to other users. The benefit of binary power control over no power control is significant when a large number of users share a few channels or more specifically when the number of users cannot be accommodated by the number of channels without causing harmful interference to each other. We found that when optimizing proportional fairness no power control (which is equivalent to binary power control since no user can have 0 capacity) causes the Jain fairness to decrease slightly, but it also causes the capacity to increase slightly. We found that for our system model that spread spectrum techniques do not tend to yield any benefit over using single-channel techniques. While a small benefit to spread spectrum techniques may exist in real scenarios, it does not in our system model because the optimization is performed over static channel gain values. Using multi-channel techniques yield better performance than single-channel techniques as the number of channels increases, but the benefit as more channels are added follows the law of diminishing returns.

We determined that the interference symmetry as defined in Section 6.1 has an impact on the benefit of collaboration. Specifically, when the channel gains are more asymmetric, it is more likely that the benefit of collaboration is higher. We compared this to a centralized network system where a distributed technique has been proven to achieve the optimal sum-rate. We determined that the benefit of channel selection increases as the number of users increases. The benefit of binary power control over no power control increases as the number of users increase and the number of channels decrease. The benefit of full power control over binary power control is negligible when optimizing sum-rate. In comparing the optimization results with sensing, we saw that fairness gains for a specific number of users and channels tend to mirror capacity gains (i.e., if the capacity gain increases as the number of users increases, the fairness gain also increases as the number of users increases). We introduced a sensing technique in which users turn off if they cannot achieve some cutoff capacity to attempt to achieve some of the gains associated with binary power control. This technique was able to bridge a significant fraction of the capacity benefit with even a small cutoff capacity value. To consider a listen-before-talk technique, we explored time division techniques in which users were favored according to the achievable capacity they could achieve. Using these type of techniques, we avoid the interference that the collaboration attempts to avoid; however,

the time division technique does not consider the possibility that two users could transmit simultaneously without causing each other harmful interference, and therefore it is unable to achieve nearly equivalent capacities to the optimal sum-rate. However, this technique has the benefit of allowing us to achieve better fairness values (at the cost of better capacity values). We estimated the overhead involved with the time division technique to be about 20% assuming a random access collision-based technique. Finally, we considered the impact of bursty traffic on the results of both the listen-before-talk and treating interference as noise techniques.

Appendix A

Derivation of SINR for Frequency Hopping SS with MRC Combining

In this section, we derive that the $SINR$ for frequency hopping spread spectrum using MRC combining is equal to the summation of the $SINR$ on the individual channels. The parameter P_r is the total received signal power. Given M channels, the fraction of power transmitted on channel k is γ_k. The channel fractions must sum to 1 or,

$$\sum_{k=1}^{M} \gamma_k = 1 \tag{A.1}$$

If the interference plus noise power for a user on channel k is denoted by I_k, the $SINR$ on channel k is,

$$SINR_k = \frac{\gamma_k \cdot P_r}{I_k} \tag{A.2}$$

If we scale the received signal on each channel k by $\sqrt{w_k}$ (or equivalently scale the power on each channel k by w_k), we get that the $SINR$ after combining is,

$$SINR = \frac{\sum_{i=1}^{M} w_i \cdot P_r \cdot \gamma_i}{\sum_{i=1}^{M} w_i \cdot I_i \cdot \gamma_i} \tag{A.3}$$

We note that I_i is scaled by γ_i because frequency hopping causes the receiver to only receive a fraction of the interference and noise energy on any particular channel. This is because frequency hopping causes the user to only spend γ_i fraction of the time on channel i (and therefore the receiver only receives γ_i fraction of the noise and interference energy on channel i). If we define,

$$w_i = \frac{1}{I_i} \tag{A.4}$$

then we have,

$$SINR = \frac{\sum_{i=1}^{M} P_r \cdot \gamma_i \cdot \frac{1}{I_i}}{\sum_{i=1}^{M} \gamma_i} = \sum_{i=1}^{M} \frac{P_r \cdot \gamma_i}{I_i} = \sum_{i=1}^{M} SINR_i \tag{A.5}$$

So using MRC causes the $SINR$ for frequency hopping spread spectrum to equal the summation of the $SINR$ on the individual channels.

Bibliography

[1] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Transactions on Information Theory*, vol. 46, pp. 388–404, Mar. 2000.

[2] J. Liu, Y. T. Hou, Y. Shi, H. D. Sherali, and S. Kompella, "On the capacity of multiuser MIMO networks with interference," *IEEE Transactions on Wireless Communications*, vol. 7, pp. 488–494, Feb. 2008.

[3] R. F. Bellaver, "AT&T in early broadcasting," in *Proc. IEEE International Symposium on Technology and Society*, Rome, Italy, Sept. 2000, pp. 293–296.

[4] M. M. Buddhikot, "Understanding dynamic spectrum access: Models, taxonomy and challenges," in *Proc. IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Dublin, Ireland, April 2007, pp. 649–663.

[5] J. Zander, "Radio resource management in future wireless networks: requirements and limitations," *IEEE Communications Magazine*, vol. 35, pp. 30–36, Aug. 1997.

[6] F. H. Sanders, "Broadband spectrum surveys in Denver, CO, San Diego, CA, and Los Angeles, CA: methodology analysis and comparative results," in *Proc. IEEE International Symposium on Electromagnetic Compatibility*, Denver, CO, Aug. 1998, pp. 988–993.

[7] R. B. Bacchus, A. J. Fertner, C. S. Hood, and D. A. Roberson, "Long-term, wide-band spectral monitoring in support of dynamic spectrum access networks at the IIT spectrum observatory," in *Proc. IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, Oct. 2008, pp. 1–10.

[8] R. Menon, A. B. MacKenzie, R. M. Buehrer, and J. H. Reed, *Accepted for Publication into IEEE Transactions on Communications*, Apr. 2009.

[9] L. Yang and G. B. Giannakis, "Ultra-wideband communications: an idea whose time has come," *IEEE Signal Processing Magazine*, vol. 21, pp. 26–54, Nov. 2004.

[10] V. Chakravarthy, Z. Wu, M. Temple, F. Garber, and X. Li, "Cognitive radio centric overlay/underlay waveform," in *Proc. IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, Oct. 2008, pp. 1–10.

[11] R. Menon, R. M. Buehrer, and J. H. Reed, "Outage probability based comparison of underlay and overlay spectrum sharing techniques," in *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Baltimore, MD, Nov. 2005, pp. 101–109.

[12] Q. Zhao and B. M. Sadler, "A survey of dynamic spectrum access," *IEEE Signal Processing Magazine*, vol. 24, pp. 79–89, May 2007.

[13] D. Niyato and E. Hossain, "Competitive pricing for spectrum sharing in cognitive radio networks: Dynamic game, inefficiency of nash equilibrium, and collusion," *IEEE Journal on Selected Areas in Communications*, vol. 26, pp. 192–202, Jan. 2008.

[14] C. Kloeck, H. Jaekel, and F. Jondral, "Auction sequence as a new resource allocation mechanism," in *IEEE 62nd Vehicular Technology Conference*, Sept. 2005, pp. 240–244.

[15] T. Kamakaris, M. M. Buddhikot, and R. Iyer, "A case for coordinated dynamic spectrum access in cellular networks," in *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Baltimore, MD, Nov. 2005, pp. 289–298.

[16] M. M. Buddhikot, P. Kolodzy, S. Miller, K. Ryan, and J. Evans, "DIMSUMNet: New directions in wireless networking using coordinated dynamic spectrum access," in *IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks*, June 2005, pp. 78–85.

[17] C. Raman, R. D. Yates, and N. B. Mandayam, "Scheduling variable rate links via a spectrum server," in *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Baltimore, MD, Nov. 2005, pp. 110–118.

[18] W. Lehr and N. Jesuale, "Spectrum pooling for next generation public safety radio systems," in *New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, Oct. 2008, pp. 1–23.

[19] W. Lehr and J. Crowcroft, "Managing shared access to a spectrum commons," in *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Baltimore, MD, Nov. 2005, pp. 420–444.

[20] A. Haghighat, "A review on essentials and technical challenges of software defined radio," in *Proceedings of MILCOM*, Montreal, Que., Canada, Oct. 2002, pp. 377–382.

[21] J. Mitola III and G. Q. Maguire Jr., "Cognitive radio: making software radios more personal," *IEEE Personal Communications*, vol. 6, pp. 13–18, Aug. 1999.

[22] J. Mitola III, "Cognitive radio for flexible mobile multimedia communications," in *Proc. IEEE International Workshop on Mobile Multimedia Communications*, San Diego, CA, Nov. 1999, pp. 3–10.

[23] C. Cordeiro, K. Challapali, and M. Ghosh, "Cognitive PHY and MAC layers for dynamic spectrum access and sharing of TV bands," in *TAPAS '06: Proceedings of the first international workshop on Technology and policy for accessing spectrum.* New York, NY, USA: ACM, 2006, p. 3.

[24] T. Yucek and H. Arslan, "A survey of spectrum sensing algorithms for cognitive radio applications," *IEEE Communications Surveys & Tutorials*, vol. 11, pp. 116–130, First-Quarter 2009.

[25] W. Zhang, R. K. Mallik, and K. B. Letaief, "Cooperative spectrum sensing optimization in cognitive radio networks," in *IEEE International Conference on Communications*, Beijing, May 2008, pp. 3411–3415.

[26] I. F. Akyildiz, W.-Y. Lee, M. C. Vuran, and S. Mohanty, "A survey on spectrum management in cognitive radio networks," *IEEE Communications Magazine*, vol. 46, pp. 40–48, Apr. 2008.

[27] P. J. Jeong and M. Yoo, "Resource-aware rendezvous algorithm for cognitive radio networks," in *International Conference on Advanced Communication Technology*, Gangwon-Do, Feb. 2007, pp. 1673–1678.

[28] A. K. Sadek, W. Zhang, and S. J. Shellhammer, "Listen-before-talk versus treating interference as noise for spectrum sharing," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access*, Chicago, IL, Oct. 2008, pp. 1–6.

[29] L. Yang, M. Kang, and M. S. Alouini, "On the capacity-fairness tradeoff in multiuser diversity systems," *IEEE Transactions on Vehicular Technology*, vol. 56, pp. 1901–1907, July 2007.

[30] C. Peng, H. Zheng, and B. Y. Zhao, "Utilization and fairness in spectrum assignment for opportunistic spectrum access," *Mobile Networks and Applications*, vol. 11, pp. 555–576, Aug. 2006.

[31] J. Huang, R. A. Berry, and M. L. Honig, "Spectrum sharing with distributed interference compensation," in *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Baltimore, MD, Nov. 2005, pp. 88–93.

[32] G. Jakimoski and K. P. Subbalakshmi, "Denial-of-service attacks on dynamic spectrum access networks," in *IEEE International Conference on Communications Workshops*, Beijing, May 2008, pp. 524–528.

[33] Y. Wu, B. Wang, and K. J. R. Liu, "Repeated spectrum sharing game with self-enforcing truth-telling mechanism," in *IEEE International Conference on Communications*, Beijing, May 2008, pp. 3583–3587.

[34] P. P. Czapski, "A survey: MAC protocols for applications of wireless sensor networks," in *IEEE Region 10 Conference*, Hong Kong, Nov. 2006, pp. 1–4.

[35] L. G. Roberts, "ALOHA packet system with and without slots and capture," *ACM SIGCOMM Computer Communication Review*, vol. 5, pp. 28–42, Apr. 1975.

[36] A. M. Glass, R. L. Brewster, and N. K. Abdulaziz, "Modelling of CSMA/CA protocol by simulation," *Electronics Letters*, vol. 24, pp. 692–694, May 1988.

[37] S. S. Lam, "Packet broadcast networks - a performance analysis of the R-ALOHA protocol," *IEEE Transactions on Computers*, vol. C-29, pp. 596–603, July 1980.

[38] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang, "MACAW: a media access protocol for wireless LAN's," in *Applications, Technologies, Architectures, and Protocols for Computer Communication*, London, United Kingdom, 1994, pp. 212–225.

[39] Y. Zhou and S. M. Nettles, "Balancing the hidden and exposed node problems with power control in CSMA/CA-based wireless networks," in *IEEE Wireless Communications and Networking Conference*, Austin, TX, Mar. 2005, pp. 683–688.

[40] G. Auer, H. Haas, and P. Omiyi, "Interference aware medium access for dynamic spectrum sharing," in *IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Dublin, Apr. 2007, pp. 399–402.

[41] J. Mo, H.-S. W. So, and J. Walrand, "Comparison of multi-channel MAC protocols," in *ACM International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, Montreal, Quebec, Canada, 2005, pp. 209–218.

[42] C. Xin, L. Ma, and C.-C. Shen, "A distributed adaptive channel assignment algorithm for dynamic spectrum access mesh networks," in *International Conference on Communications and Networking*, Hangzhou, China, Aug. 2008, pp. 1178–1182.

[43] P. Bahl, R. Chandra, and J. Dunagan, "SSCH: slotted seeded channel hopping for capacity improvement in IEEE 802.11 ad-hoc wireless networks," in *Proceedings of the 10th Annual international Conference on Mobile Computing and Networking*, Philadelphia, PA, Sept. 2004, pp. 216–230.

[44] P. Piggin and K. L. Stanwood, "Standardizing WiMAX solutions for coexistence in the 3.65 ghz band," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, Oct. 2008, pp. 1–7.

[45] C. Xin, B. Xie, and C.-C. Shen, "A novel layered graph model for topology formation and routing in dynamic spectrum access networks," in *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Baltimore, MD, Nov. 2005, pp. 308–317.

[46] Y. Xing, R. Chandramouli, S. Mangold, and S. S. N, "Dynamic spectrum access in open spectrum wireless networks," *IEEE Journal on Selected Areas in Communications*, vol. 24, pp. 626–637, Mar. 2006.

[47] D. T. C. Wong, A. T. Hoang, Y.-C. Liang, and F. P. S. Chin, "Dynamic spectrum access with virtual partitioning in open spectrum wireless networks," in *IEEE Vehicular Technology Conference*, Singapore, May 2008, pp. 1584–1588.

[48] D. T. C. Wong, A. T. Hoang, Y.-C. Liang, and F. Chin, "Dynamic spectrum access with prioritization in open spectrum wireless networks," in *IEEE Singapore International Conference on Communication Systems*, Guangzhou, China, Nov. 2008, pp. 1026–1030.

[49] B. Wang, Z. Ji, and K. J. R. Liu, "Primary-prioritized markov approach for dynamic spectrum access," in *IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Dublin, Apr. 2007, pp. 507–515.

[50] ——, "Self-learning repeated game framework for distributed primary prioritized dynamic spectrum access," in *IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks*, San Diego, CA, June 2007, pp. 631–638.

[51] M. Raspopovic and C. Thompson, "Finite population model for performance evaluation between narrowband and wideband users in the shared radio spectrum," in *IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Dublin, Apr. 2007, pp. 340–346.

[52] L. Cao and H. Zheng, "Distributed rule-regulated spectrum sharing," *IEEE Journal on Selected Areas in Communications*, vol. 26, pp. 130–145, Dec. 2008.

[53] H. Zheng and L. Cao, "Device-centric spectrum management," in *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Baltimore, MD, Nov. 2005, pp. 56–65.

[54] M. E. Steenstrup, "Opportunistic use of radio-frequency spectrum: A network perspective," in *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Baltimore, MD, Nov. 2005, pp. 638–641.

[55] Y. Shi and Y. T. Hou, "Optimal power control for multi-hop software defined radio networks," in *IEEE INFOCOM*, Anchorage, AK, May 2007, pp. 1694–1702.

[56] ——, "A distributed optimization algorithm for multi-hop cognitive radio networks," in *IEEE INFOCOM*, Phoenix, AZ, Apr. 2008, pp. 1292–1300.

[57] M. Y. E. Nainay, D. H. Friend, and A. B. MacKenzie, "Channel allocation and power control for dynamic spectrum cognitive networks using a localized island genetic algorithm," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, Oct. 2008, pp. 1–5.

[58] D. H. Friend, M. Y. ElNainay, Y. Shi, and A. B. MacKenzie, "Architecture and performance of an island genetic algorithm-based cognitive network," in *IEEE Consumer Communications and Networking Conference*, Las Vegas, NV, Jan. 2008, pp. 993–997.

[59] P. Yang and G. Chen, "FAST CASH: Fair and stable channel assignment on heterogeneous wireless mesh network," in *International Conference for Young Computer Scientists*, Hunan, China, Nov. 2008, pp. 451–456.

[60] F. Berggren, "Capacity analysis for noncooperative interference environments," in *IEEE 61st Vehicular Technology Conference*, Beijing, May 2005, pp. 1479–1483.

[61] G. Alyfantis, G. Marias, S. Hadjiefthymiades, and L. Merakos, "Non-cooperative dynamic spectrum access for CDMA networks," in *IEEE Global Telecommunications Conference*, Washington D.C., Nov. 2007, pp. 3574–3578.

[62] D. Malone, P. Clifford, D. Reid, and D. J. Leith, "Experimental implementation of optimal WLAN channel selection without communication," in *IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Dublin, Apr. 2007, pp. 316–319.

[63] D. J. Leith and P. Clifford, "A self-managed distributed channel selection algorithm for WLANs," in *International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, Apr. 2006, pp. 1–9.

[64] M. Sharma, A. Sahoo, and K. D. Nayak, "Channel selection under interference temperature model in multi-hop cognitive mesh networks," in *IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Dublin, Apr. 2007, pp. 133–136.

[65] B. Babadi and V. Tarokh, "The impact of spectrum sensing time on the performance of the GADIA algorithm," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, Oct. 2008, pp. 1–7.

[66] H. Qin, H. Wang, and H. Zhou, "A selfish game-theoretic approach for cognitive radio networks with dynamic spectrum sharing," in *International Conference on Computer Science and Software Engineering*, Wuhan, China, Dec. 2008, pp. 1105–1109.

[67] S. Bahramian and B. H. Khalaj, "Joint dynamic frequency selection and power control for cognitive radio networks," in *International Conference on Telecommunications*, St. Petersburg, June 2008, pp. 1–6.

[68] D. Li, X. Dai, and H. Zhang, "Game theoretic analysis of joint rate and power control in cognitive radio networks," in *International Conference on Communications, Circuits and Systems*, Fujan, May 2008, pp. 319–322.

[69] Y. Xing and R. Chandramouli, "QoS constrained secondary spectrum sharing," in *First IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Baltimore, MD, Nov. 2005, pp. 658–661.

[70] Y. Xing, C. N. Mathur, M. A. Haleem, R. Chandramouli, and K. P. Subbalakshmi, "Dynamic spectrum access with QoS and interference temperature constraints," *IEEE Transactions on Mobile Computing*, vol. 6, pp. 423–433, Apr. 2007.

[71] Y. Xing, C. N. Mathur, M. A. Haleem, R. Chadramouli, and K. P. Subbalakshmi, "Priority based dynamic spectrum access with QoS and interference temperature constraints," in *IEEE International Conference on Communications*, Istanbul, June 2006, pp. 4420–4425.

[72] D. I. Kim, L. B. Le, and E. Hossain, "Joint rate and power allocation for cognitive radios in dyanmic spectrum access environment," *IEEE Transactions on Wireless Communications*, vol. 7, pp. 5517–5527, Dec. 2008.

[73] B. Wang, Y. Wu, Z. Ji, K. J. Liu, and T. Clancy, "Game theoretical mechanism design methods," *IEEE Signal Processing Magazine*, vol. 25, pp. 74–84, Nov. 2008.

[74] D. Xu, E. Jung, and X. Liu, "Optimal bandwidth selection in multi-channel cognitive radio networks," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, Oct. 2008, pp. 1–11.

[75] N. Nie and C. Comaniciu, "Adaptive channel allocation spectrum etiquette for cognitive radio networks," in *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Baltimore, MD, Nov. 2005, pp. 269–278.

[76] J. Huang, R. A. Berry, and M. L. Honig, "Distributed interference compensation for wireless networks," *IEEE Journal on Selected Areas in Communications*, vol. 24, pp. 1074–1084, May 2006.

[77] R. Etkin, A. Parekh, and D. Tse, "Spectrum sharing for unlicensed bands," *IEEE Journal on Selected Areas in Communications*, vol. 25, pp. 517–528, Apr. 2007.

[78] Z. Ji and K. J. R. Liu, "Dynamic spectrum sharing: A game theoretical overview," *IEEE Communications Magazine*, vol. 45, pp. 88–94, May 2007.

[79] N. Clemens and C. Rose, "Intelligent power allocation strategies in an unlicensed spectrum," in *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Baltimore, MD, Nov. 2005, pp. 37–42.

[80] S. Subramani, T. Basar, S. Armour, D. Kaleshi, and Z. Fan, "Noncooperative equilibrium solutions for spectrum access in distributed cognitive radio networks," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, Oct. 2008, pp. 1–5.

[81] C. W. Sung and K. K. Leung, "A framework for opportunistic power control - convergence and applications," *IEEE Transactions on Information Theory*, vol. 51, pp. 2625–2635, July 2005.

[82] G. J. Foschini and Z. Miljanic, "A simple distributed autonomous power control algorithm and its convergence," *IEEE Transactions on Vehicular Technology*, vol. 42, pp. 641–646, Nov. 1993.

[83] S. A. Grandhi, R. Vijayan, D. J. Goodman, and J. Zander, "Centralized power control in cellular radio systems," *IEEE Transactions on Vehicular Technology*, vol. 42, pp. 466–468, Nov. 1993.

[84] S. Hanly and D. N. C. Tse, "Multiaccess fading channels-part i: Polymatroid structure, optimal resource allocation and throughput capacities," *IEEE Transactions on Information Theory*, vol. 44, pp. 2796–2815, Nov. 1998.

[85] S. Chan, "Shared spectrum access for the DoD," in *IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Dublin, Ireland, Apr. 2007, pp. 524–534.

[86] P. Pawelczak, R. V. Prasad, L. Xia, and I. G. M. M. Niemegeers, "Cognitive radio emergency networks: Requirements and design," in *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Baltimore, MD, Nov. 2005, pp. 601–606.

[87] N. Jesuale and B. C. Eydt, "A policy proposal to enable cognitive radio for public safety and industry in the land mobile radio bands," in *IEEE Internaional Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Dublin, Ireland, Apr. 2007, pp. 66–77.

[88] D. Cavalcanti and M. Ghosh, "Cognitive radio networks: Enabling new wireless broadband opportunities," in *3rd International Conference on Cognitive Radio Oriented Wireless Networks and Communications*, Singapore, May 2008, pp. 1–6.

[89] M. Sherman, A. N. Mody, R. Martinez, C. Rodriguez, and R. Reddy, "IEEE standards supporting cognitive radio and networks, dynamic spectrum access and coexistence," *IEEE Communications Magazine*, vol. 46, pp. 72–79, July 2008.

[90] J. B. MacQueen, "Some methods for classification and analysis of multivariate observations," in *Berkeley Symposium on Mathematical Statistics and Probability*, 1967, pp. 281–297.

[91] F. P. Kelly, "Charging and rate control for elastic traffic," *European Transactions on Telecommunications*, vol. 8, pp. 33–37, Jan. 1997.

[92] W. Wang, T. Peng, and W. Wang, "Optimal power control under interference temperature constraints in cognitive radio network," in *IEEE Wireless Communications and Networking Conference*, Kowloon, Mar. 2007, pp. 116–120.

[93] H. D. Sherali and W. P. Adams, *A Reformulation-Linearization-Technique for Solving Discrete and Continuous Nonconvex Problems*. Boston, MA: Kluwer Academic Publishing, 1999.

[94] R. M. Buehrer and R. Mahajan, "On the usefulness of outer-loop power control with successive interference cancellation," *IEEE Transactions on Communications*, vol. 51, pp. 2091–2102, Dec. 2003.

[95] J. Zander, "Performance of optimum transmitter power control in cellular radio systems," *IEEE Transactions on Vehicular Technology*, vol. 41, pp. 57–62, Feb. 1992.

[96] A. Gjendemsi, D. Gesbert, G. E. Oien, and S. G. Kiani, "Binary power control for sum rate maximization over multiple interfering links," *IEEE Transactions on Wireless Communications*, vol. 7, pp. 3164–3173, Aug. 2008.

[97] J. Weinmiller, M. Schlager, A. Festag, and A. Wolisz, "Performance study of access control in wireless LANs - IEEE 802.11 DFWMAC and ETSI RES 10 hiperlan," *Mobile Networks and Applications*, vol. 2, pp. 55–67, 1997.

[98] R. Etkin, D. N. C. Tse, and H. Wang, "Gaussian interference channel capacity to within one bit: the symmetric case," in *IEEE Information Theory Workshop*, Chengdu, Oct. 2006, pp. 601–605.

[99] N. Badruddin, S. R. Bhaskaran, J. Evans, and S. Hanly, "Maximizing the sum rate in symmetric networks of interfering links under flat power constraints," in *Allerton Conference on Communication, Control, and Computing*, Urbana-Champaign, IL, Sept. 2008, pp. 46–53.

[100] O. Kaya and S. Ulukus, "Achieving the capacity region boundary of fading CDMA channels via generalized iterative waterfilling," *IEEE Transactions on Wireless Communications*, vol. 5, pp. 3215–3223, Nov. 2006.

[101] S. G. Kiani and D. Gesbert, "Capacity maximizing power allocation for interfering wireless links: A distributed approach," in *IEEE Global Telecommunications Conference*, Washington, D.C., Nov. 2007, pp. 1405–1409.

[102] S. Ulukus and R. D. Yates, "Iterative construction of optimum signature sequence sets in synchronous CDMA systems," *IEEE Transactions on Information Theory*, vol. 47, pp. 1989–1998, July 2001.

[103] P. Anigstein and V. Anantharam, "Ensuring convergence of the MMSE iteration for interference avoidance to the global optimum," *IEEE Transactions on Information Theory*, vol. 49, pp. 873–885, Apr. 2003.

[104] O. Popescu, "Interference avoidance for wireless systems with multiple receivers," Ph.D. dissertation, Rutgers University, 2004.

[105] A. Gupta, X. Lin, and R. Srikant, "Low-complexity distributed scheduling algorithms for wireless networks," in *IEEE International Conference on Computer Communications*, Anchorage, AK, May 2007, pp. 1631–1639.

[106] A. Giridhar and P. R. Kumar, "Distributed clock synchronization over wireless networks: Algorithms and analysis," in *IEEE Conference on Decision and Control*, San Diego, CA, Dec. 2006, pp. 4915–4920.

[107] L. Kleinrock and F. Tobagi, "Packet switching in radio channels: Part 1 - carrier sense multiple-access models and their throughput delay characteristics," *IEEE Transactions on Communications*, vol. 23, pp. 1400–1416, Dec. 1975.

[108] N. Abramson, "The ALOHA system - another alternative for computer communications," in *AFIPS '70 (Fall): Proceedings of the November 17-19, 1970, fall joint computer conference*, Houston, Texas, Nov. 1970, pp. 281–285.

[109] V. Bharghavan, "Performance evaluation of algorithms for wireless medium access," in *IEEE International Computer Performance and Dependability Symposium*, Durham, NC, Sept. 1998, pp. 86–95.